An update on *Might*

Jaap van der Does, Willem Groeneveld and Frank Veltman

ILLC, University of Amsterdam

1 Introduction

The last decade has shown a growing concern with aspects of language interpretation for which the truth conditional paradigm seems too poor. The focus is now on the context change potential of a sentence rather than on its truth conditions. Prime example is the dynamic semantics for predicate logic introduced in [Groenendijk and Stokhof 1991].

Besides introducing new techniques in formal semantics, these dynamic systems also offer new challenges for logic, because they allow for more variation in the notion of consequence. These new notions of consequence differ from the classical 'preservation of truth' notion in various respects. They have different structural properties, which means, amongst other things, that standard constructions for proving completeness do not apply.

In this paper we will be concerned with the semantics for 'might' and the three consequence relations in the update semantics of [Veltman 1991]. The consequence relations are introduced and studied in an abstract setting in section 2. Next we will turn to the logics for 'might' that these consequence relations give rise to. We will present three sequent-style systems; each of them is shown to be complete and decidable. The final section contains a general cut elimination result.[1]

[1]This paper grew out of a talk by Frank Veltman in which he presented three complete sequent systems for 'might' (corresponding to the three notions of validity). Jaap van der Does streamlined the sequent systems so that they lend themselves to (Cautious) Cut elimination [van der Does 1994], and Willem Groeneveld supplied in his thesis the more general outlook that is now presented in section 2. See [Groeneveld and Veltman 1994] (section 3). Cf. also [Groeneveld 1995]. The first author takes part in the PIONIER project 'Reasoning with Uncertainty' (NWO grant PGS –22–262). For the other two authors the work for this paper formed part of the Esprit Basic Research Project DYANA (6852). Some of the results presented here are also proved by [Van Eijck and De Vries 1995] by means of a Hoare logic, and a translation into S5.

2 Abstract Update Semantics

A system of dynamic semantics for a given language specifies a set of information states and for each sentence a binary relation between information states. In this section we completely abstract away from any specific feature of the information states, and from any syntactic detail of the language under consideration.

Definition 2.1 (Frames) Suppose $\mathcal{L}$ is some non-empty set of symbols. Then a *frame* for $\mathcal{L}$ is a structure $\mathcal{F} = (\Sigma, ([\varphi]_{\mathcal{F}})_{\varphi \in \mathcal{L}})$, where Σ is a non-empty set, and for each $\varphi \in \mathcal{L}$, $[\varphi]_{\mathcal{F}} \subseteq \Sigma \times \Sigma$ is a binary relation on Σ. A frame $\mathcal{F}$ is *generated* if there is a unique state $\top \in \Sigma$, called the minimal state, such that for each $\sigma \in \Sigma$, $\sigma = \top$ or there are $\varphi_1, \ldots, \varphi_n \in \mathcal{L}$ such that $(\top, \sigma) \in [\varphi_1]_{\mathcal{F}} \circ \cdots \circ [\varphi_n]_{\mathcal{F}}$, where $\circ$ is relational composition. $\Box$

The three consequence relations introduced in the more concrete semantics of [Veltman 1991] can already be defined in the present abstract setting.

Definition 2.2 (Consequence relations) Let $\mathcal{L}$ be a set and consider a frame $\mathcal{F} = (\Sigma, ([p]_{\mathcal{F}})_{p \in \mathcal{L}})$. We define three concepts of validity for an argument $p_1, \ldots, p_n \Rightarrow q$, i.e. an argument with a sequence of premisses $p_1, \ldots, p_n$ and a conclusion q. [2]

1. Test Consequence:
   $\mathcal{F} \models_{\mathrm{tc}} p_1, \ldots, p_n \Rightarrow q$ iff $\mathit{fix}([p_1]_{\mathcal{F}}) \cap \cdots \cap \mathit{fix}([p_n]_{\mathcal{F}}) \subseteq \mathit{fix}([q]_{\mathcal{F}})$
2. Ignorant Consequence:
   $\mathcal{F} \models_{\mathrm{ic}} p_1, \ldots, p_n \Rightarrow q$ iff $(\top, t) \in [p_1]_{\mathcal{F}} \circ \cdots \circ [p_n]_{\mathcal{F}}$ implies $(t, t) \in [q]_{\mathcal{F}}$
3. Update Test Consequence:
   $\mathcal{F} \models_{\mathrm{utc}} p_1, \ldots, p_n \Rightarrow q$ iff $\mathit{rge}([p_1]_{\mathcal{F}} \circ \cdots \circ [p_n]_{\mathcal{F}}) \subseteq \mathit{fix}([q]_{\mathcal{F}})$ $\Box$[3]

When $(i, i) \in [\varphi]$, i is a fixed point of the relation $[\varphi]$. This means that an update with φ does not force an agent who happens to be in state i to change it. Apparently, the information supplied by φ is already contained in i: 'the state i *supports* φ', as we shall say, or 'φ is *accepted* in i'.

[2] *fix*, *do* and *rge* respectively stand for the operation of taking the set of fixed points, the domain, and the range of a binary relation.

[3] The subscripts are mnemonic: 'tc' stands for 'test consequence', 'ic' for ignorant consequence, and 'utc' for update test consequence.

Given this, it is easy to see that the consequence relations of definition 2.2 are all based on the idea that an argument is valid if and only its conclusion is supported by its premises.

In $\models_{\text{tc}}$ this idea is made more precise by requiring that the conclusion be supported by any state which supports the premises. In the relations $\models_{\text{ic}}$ and $\models_{\text{utc}}$ a more dynamic interpretation is given. According to $\models_{\text{utc}}$ an argument is valid iff its conclusion is supported by any state in which one can arrive after learning the premisses (in the order in which they are given). The relation $\models_{\text{ic}}$ is a special case of this: it says that an argument is valid if starting from the state of ignorance any update with the premisses will lead to a state which supports the conclusion.

In general $\models_{\text{utc}}$ implies $\models_{\text{tc}}$ and $\models_{\text{ic}}$, as is easily verified. The following proposition shows that the converse implications also hold provided the following constraints are satisfied:

- **Permutation:** for all $p, q \in \mathcal{L}$, $[p]_{\mathcal{F}} \circ [q]_{\mathcal{F}} = [q]_{\mathcal{F}} \circ [p]_{\mathcal{F}}$
- **Idempotency:** for all $p \in \mathcal{L}$, $rge([p]_{\mathcal{F}}) \subseteq fix([p]_{\mathcal{F}})$

Proposition 2.3 Let $\mathcal{F} = (\Sigma, ([p]_{\mathcal{F}})_{p \in \mathcal{L}})$ be a frame for $\mathcal{L}$. Then the following hold.

1. If $\mathcal{F} \models_{\text{utc}} p_1, \ldots, p_n \Rightarrow q$ then $\mathcal{F} \models_{\text{tc}} p_1, \ldots, p_n \Rightarrow q$
2. If $\mathcal{F} \models_{\text{utc}} p_1, \ldots, p_n \Rightarrow q$ then $\mathcal{F} \models_{\text{ic}} p_1, \ldots, p_n \Rightarrow q$.
3. If $\mathcal{F}$ satisfies Permutation and Idempotency, then $\mathcal{F} \models_{\text{tc}} p_1, \ldots, p_n \Rightarrow q$ implies $\mathcal{F} \models_{\text{utc}} p_1, \ldots, p_n \Rightarrow q$
4. If $\mathcal{F}$ is generated, and satisfies Permutation and Idempotency, then $\mathcal{F} \models_{\text{ic}} p_1, \ldots, p_n \Rightarrow q$ implies $\mathcal{F} \models_{\text{utc}} p_1, \ldots, p_n \Rightarrow q$.

Proof: left to the reader. □

Given that the three consequence relations do not in general coincide, the question arises in which respects they differ. We will approach this question by proving that the relations validate different structural inference rules.

Definition 2.4 (Pure Structural Rule) Let $d \in \{cl, icl, utc\}$. Define $Val(d, \mathcal{F}) = \{(X \Rightarrow p) \mid \mathcal{F} \models_d (X \Rightarrow p)\}$. Then a sequent rule R is a *pure structural rule* for d provided that for all frames $\mathcal{F}$, $Val(d, \mathcal{F})$ is closed under R. □

For example, a simple inspection of the definition of $\models_{\text{utc}}$ will show that the rules of Left Monotony and Cautious Cut

$$\frac{X \Rightarrow q}{pX \Rightarrow q}\ LM \qquad\qquad \frac{X \Rightarrow p \quad XpY \Rightarrow q}{XY \Rightarrow q}\ CC$$

are pure structural rules of $\models_{\text{utc}}$, and in fact and any rule that is derivable from LM and CC will also be a pure structural rule for $\models_{\text{utc}}$. A natural question is then whether the rules LM and CC completely determine the pure structural rules of $\models_{\text{utc}}$. The answer to this question is yes: if R is a pure structural rule of $\models_{\text{utc}}$ then it is derivable from LM and CC.

Definition 2.5 (Structural Completeness) A set of sequent rules Σ is structurally complete for d if every pure structural rule of d is derivable from Σ ($d \in \{cl, icl, utc\}$). □

The structural completeness results presented below are all based on representation results in the following sense.

Definition 2.6 A set of sequent rules Σ is d-representable if and only if for each set of sequents Δ that is closed under all rules of Σ, there exists a frame $\mathcal{F}_d(\Delta)$ such that: $\mathcal{F}_d(\Delta) \models_d (X \Rightarrow p)$ iff $(X \Rightarrow p) \in \Delta$. □

Lemma 2.7 If Σ is d-representable then Σ is structurally complete for d.
Proof: Suppose that Σ is d-representable, but there is some pure structural rule $S_1, \ldots, S_n/S$ for d that is not derivable from Σ. Let Δ be the closure of $\{S_1, \ldots, S_n\}$ under all rules of Σ. Then $S \notin \Delta$. Then $\mathcal{F}_d(\Delta) \models_d S_1$, $\ldots$, $\mathcal{F}_d(\Delta) \models_d S_n$, but $\mathcal{F}_d(\Delta) \not\models_d S$. So $Val(d, \mathcal{F}_d(\Delta))$ is not closed under $S_1, \ldots, S_n/S$, contradiction. □

Consider these structural rules of inference:

$$\frac{}{p \Rightarrow p}\ Refl \qquad \frac{XY \Rightarrow q}{XpY \Rightarrow q}\ Mon \qquad \frac{X \Rightarrow p \quad YpZ \Rightarrow q}{YXZ \Rightarrow q}\ Cut$$

$$\frac{XpYpZ \Rightarrow q}{XpYZ \Rightarrow q}\ Contr \qquad \frac{X \Rightarrow q}{pX \Rightarrow q}\ LM \qquad \frac{X \Rightarrow p \quad XpY \Rightarrow q}{XY \Rightarrow q}\ CC$$

The rules $Refl$ (Reflexivity), Cut, Mon (Monotony) and $Contr$ (Contraction) are all familiar from classical logic. In fact any logic that defines valid consequence as preservation of truth will validate these four structural rules. LM (Left Monotony) is a weak variant of Monotony, and CC (Cautious Cut) is a weak variant of Cut.

Proposition 2.8 On the class of all frames:

(i) *Refl*, *Cut*, *Mon* and *Contr* are structurally complete for $\models_{\text{tc}}$.

(ii) *CC* completely determines the structural rules of $\models_{\text{ic}}$.

(iii) *CC* and *LM* are structurally complete for $\models_{\text{utc}}$.

Proof: the details for (i) can be found in [van Benthem 1991a] or [Groeneveld 1995]. The main idea is to use lemma 2.7, and to represent an any set of sequents Δ that is closed under Reflexivity, Cut, Monotony and Contraction, by a frame $Cl(\Delta) = (\Sigma, ([p])_{p\in\mathcal{L}})$, which is defined by $\Sigma = \mathcal{L}^{<\omega}$ and $[p] = \{(X,X) \mid (X \Rightarrow p) \in \Delta\}$.

The proof of (ii) is more simple than (iii); for the latter we cite the proof of [van Benthem 1991c]. Suppose Δ is closed under LM and CC, and define a frame by $(\Sigma, ([p])_{p\in\mathcal{L}})$ by $\Sigma = \mathcal{L}^{<\omega}$, and

$$[p] = \{(X,X) \mid (X \Rightarrow p) \in \Delta\} \cup \{(X, Xp) \mid X \in \mathcal{L}^{<\omega}\}$$

Claim: $(X \Rightarrow p) \in \Delta$ iff $Utc(\Delta) \models_{\text{utc}} X \Rightarrow p$. This means that Δ is $\models_{\text{utc}}$-representable, which is sufficient in view of lemma 2.7.

Proof of the claim: Assume that $(p_1, \ldots, p_n \Rightarrow q) \in \Delta$ and suppose that $X \in rge([p_1] \circ \cdots \circ [p_n])$. Then $X_0[p_1]X_1 \cdots X_{n-1}[p_n]X_n = X$ for some $X_0, \ldots, X_n$. We show with induction on i that $(X_i p_{i+1}, \ldots, p_n \Rightarrow q) \in \Delta$. By taking $i = n$, this gives that $(X \Rightarrow q) \in \Delta$, so $X \in fix[q]$, which we are after.

$i = 0$: The assumption that $(p_1, \ldots, p_n \Rightarrow q) \in \Delta$ and LM yields that $(X_0 p_1, \ldots, p_n \Rightarrow q) \in \Delta$.

$i \rightarrow i+1$: suppose as induction hypothesis that $(X_i p_{i+1}, \ldots, p_n \Rightarrow q) \in \Delta$. We have $X_i[p_{i+1}]X_{i+1}$. By the definition of $[\cdot]$, one of the following two cases must obtain. Case 1: $X_{i+1} = X_i p_{i+1}$. The apply the induction hypothesis. Case 2: $X_{i+1} = X_i$ and $(X_i \Rightarrow p_{i+1}) \in \Delta$. Then apply CC, i.e from $X_i \Rightarrow p_{i+1}$ and $X_i p_{i+1}, \ldots, p_n \Rightarrow q$ we conclude $X_i p_{i+2}, \ldots, p_n \Rightarrow q$.

Conversely, suppose that $rge([p_1] \circ \cdots \circ [p_n]) \subseteq fix([q])$. Since $p_1 \ldots p_n \in rge([p_1] \circ \cdots \circ [p_n])$, also $p_1 \ldots p_n \in fix([q])$. But the latter can only be if $(p_1, \ldots, p_n \Rightarrow q) \in \Delta$. □

The results of this proposition, for instance for $\models_{\text{utc}}$, should be seen in the following perspective. The content of proposition 2.8 (ii) is that for any concrete relational semantics with $\models_{\text{utc}}$ as consequence relation, LM and CC will be valid structural inference rules. This means that any additional structural inference rule that comes out as valid in such a concrete semantics does not so much reflect a property of $\models_{\mathrm{utc}}$, but rather reflects a property that is specific for the semantics.

An instance of this phenomenon can be seen by restricting the class of all relational frames to the functional and idempotent frames (the 'might' semantics we will discuss below has these two properties).

$$\frac{}{Xp \Rightarrow p}\ RRefl \qquad\qquad \frac{X \Rightarrow p \quad XY \Rightarrow q}{XpY \Rightarrow q}\ CM$$

Proposition 2.9 On the class of all functional idempotent frames:

(i) *Refl*, *Cut*, *Mon* and *Contr* are structurally complete for $\models_{\mathrm{tc}}$.

(ii) *CC*, *RRefl* and *CM* are structurally complete for $\models_{\mathrm{ic}}$.

(iii) *CC*, *LM* , *RRefl* and *CM* are structurally complete for $\models_{\mathrm{utc}}$.

Proof: For (i), note that the frame used as the relational representation for $\models_{\mathrm{tc}}$ in proposition 2.8 is in fact idempotent and *partial* functional. This can be turned into a functional frame by the familiar trick of adding a new state $\perp$ to which all missing transitions are directed.

Again $\models_{\mathrm{ic}}$ is simpler than $\models_{\mathrm{utc}}$ so only the latter is treated here. One can adapt the relational representation for $\models_{\mathrm{utc}}$ to

$$[p] = \{(X,X) \mid (X \Rightarrow p) \in \Gamma\} \cup \{(X,Xp) \mid (X \Rightarrow p) \notin \Gamma\}$$

This yields a functional and idempotent model. The details, which are similar to the relational case, are left to the reader. $\square$

So, for $\models_{\mathrm{tc}}$ nothing changes, but the dynamic relations $\models_{\mathrm{ic}}$ and $\models_{\mathrm{utc}}$ get two extra structural rules, *RRefl* and *CM*. In fact, for $\models_{\mathrm{utc}}$ the correspondence is that *RRefl* characterizes idempotency, and *CM* is valid due to functionality. For a detailed discussion of $\models_{\mathrm{utc}}$ in the context of idempotent and functional updates the reader is referred to [Groeneveld 1995].

3 'Might': syntax and semantics

In this section the syntax and semantics of the propositional language **M** is presented.

The syntax of **M** is defined on top of a standard propositional language over a set of propositional letters $\mathcal{P} := \{\mathrm{p}_1, \ldots, \mathrm{p}_{\mathrm{n}}, \ldots\}$. There is a distinction between $\mathcal{L}_0$ and $\mathcal{L}_1$ formulas in order to preclude iterations of the might-operator.

Definition 3.1 (syntax) $\mathcal{L}_0 \quad F_0 ::= p_{\rm i} \mid \neg F_0 \mid (F_0 \wedge F_0)$
$\mathcal{L}_1 \quad F_1 ::= F_0 \mid \mathrm{M} F_0$ □

Semantically, MF (read 'It might the case that F'), is interpreted as an operator which tests for consistency, a metaproperty. So, the fact that M only occurs as an outermost operator corresponds to a strict division between the object language $\mathcal{L}_0$ and the metalanguage $\mathcal{L}_1$. Below, formulas of the form Mφ are called M-formulas.

The semantics of **M** specifies the update function associated with a formula, not its truth conditions. More precisely, a formula φ denotes a function from information states to information states.

Definition 3.2 (information structures, models, updates) An *information structure* is a structure $\mathcal{I} := \langle \mathrm{I},^{\mathrm{c}}, \wedge, \top, \mathrm{F}_i \rangle$ consisting of a Boolean algebra (hence: BA) $\langle \mathrm{I},^{\mathrm{c}}, \wedge, \top \rangle$ with a family of operators $\mathrm{F}_i : \mathrm{I}^n \longrightarrow \mathrm{I}$.[4] A *model* for a vocabulary $\mathcal{P}$ is an information structure $\mathcal{I} := \langle \mathrm{I},^{\mathrm{c}}, \wedge, \top, [\![p]\!]^{\mathcal{I}} \rangle_{p \in \mathcal{P}}$, where the *update functions* $[\![p]\!]^{\mathcal{I}} : \mathrm{I} \longrightarrow \mathrm{I}$ must satisfy:

a.	$i[\![p]\!] \subseteq i$	introspective
b.	If $i \subseteq j$ then $i[\![p]\!] \subseteq j[\![p]\!]$	monotone
c.	If $i \subseteq j[\![p]\!]$ then $i \subseteq i[\![p]\!]$	stable

Here $\subseteq$ is defined by $i \subseteq j$ iff $i \wedge j = j$, as usual. When no confusion is likely the superscript $^{\mathcal{I}}$ is omitted. The argument is placed before the function so as not to disturb the order among the formulas in case of the sequences introduced below.

With every formula φ an update function $[\varphi]^{\mathcal{I}} : \mathrm{I} \longrightarrow \mathrm{I}$ is given as follows.

$$
\begin{array}{llcl}
a. & i[\mathrm{p}] & = & i[\![p]\!] \\
b. & i[\neg\varphi] & = & i - i[\varphi] \\
c. & i[\varphi \wedge \psi] & = & i[\varphi] \wedge i[\psi] \\
d. & i[\mathrm{M}\varphi] & = & \begin{cases} i & \text{if } i[\varphi] \neq \bot \\ \bot & \text{otherwise} \end{cases}
\end{array}
$$

Here $\bot =_{df} \top^c$ is the inconsistent information state, and $i - j =_{df} i \wedge j^c$. Finally, the update function of a sequence of formulas $\sigma_1 \ldots \sigma_n$ is defined inductively by $i[\sigma_1 \ldots \sigma_n]^{\mathcal{I}} = (i[\sigma_1 \ldots \sigma_{n-1}]^{\mathcal{I}})[\sigma_n]^{\mathcal{I}}$. □

[4] Cf. [Van Benthem 1991c] and [Kanazawa 1994a,b] for a more general relational notion of information structure. We will use basic facts concerning BA's without much notice.

Perhaps the most prominent feature of definition 3.2 is that all formulas are interpreted as operators on a BA. This move to a higher level enables a uniform definition of interpretation. But $\mathcal{L}_0$ formulas can also be interpreted as elements of a BA.

Proposition 3.3 For all $\varphi \in \mathcal{L}_0$ and all states i: $i[\varphi] = i \wedge \top[\varphi]$.
Proof: Induction on the structure of φ. The atomic case is based on the fact that $[p]$ is introspective, monotone, and stable. $\square$

Proposition 3.3 reveals to what extent the constraints on $[\![p]\!]$ are preserved under the definition of $[\varphi]$:

Corollary 3.4 For all formulas φ, $[\varphi]$ is introspective and monotone. It is stable in case φ is a formula of $\mathcal{L}_0$. $\square$

A further consequence of 3.3 is that $[\varphi]$ is idempotent: $i[\varphi][\varphi] = i[\varphi]$, for all φ and all $i \in \mathcal{I}$. Proposition 3.3 does not hold for M-formulas. The next proposition collects some useful properties of sequences of formulas.

Proposition 3.5 For all sequences Π, Π', and all $i \in \mathcal{I}$:

i) $\perp[\Pi]^{\mathcal{I}} = \perp$

ii) $i[\Pi, \Pi']^{\mathcal{I}} = (i[\Pi]^{\mathcal{I}})[\Pi']^{\mathcal{I}}$

iii) $i[\Pi, \sigma, \Pi'] \subseteq i[\Pi, \Pi']$

iv) $i[\Pi] \subseteq i[\Pi_0]$

v) $i[\Pi] \neq \perp$ iff $i[\Pi^*] = i[\Pi^*_0] \neq \perp$ for each initial segment Π^* of Π.

Π_0 is Π with all M-formulas erased.
Proof: We only prove (v). One direction is clear, so assume $i[\Pi] \neq \perp$. First note that no initial segment Π^* has $i[\Pi^*] = \perp$; otherwise, $i[\Pi] = \perp$ by (i–ii). As to the remaining claim, distinguish two cases in an induction on the length of Π^*.

- $\Pi^* \equiv \Lambda, \varphi$ with $\varphi \in \mathcal{L}_0$. Then: $i[\Lambda, \varphi] =_{\text{i.h.}} i[\Lambda_0, \varphi] = i[(\Lambda, \varphi)_0]$.
- $\Pi^* \equiv \Lambda, \mathrm{M}\varphi$. It is an immediate consequence of $i[\Lambda, \varphi] = \perp$ that $i[\Pi^*] = \perp$, which is impossible. So, $i[\Lambda, \varphi] \neq \perp$, and therefore $i[\Lambda, \mathrm{M}\varphi] = i[\Lambda] =_{\text{i.h.}} i[\Lambda_0] = i[(\Lambda, \mathrm{M}\varphi)_0]$. $\square$

The models introduced in [Veltman 1991] are more concrete than the ones defined above. They are based on the following three assumptions. First, a world is a finite set of proposition letters, which represent the atomic facts that obtain in it. Second, an information state is a set of worlds, the worlds compatible with the information at hand. Third, a model should contain as many information states as possible. More in particular, given a finite set of proposition letters $\mathcal{P}$, the models have the form $\langle\wp\wp(\mathcal{P}),^c,\cap,\wp(\mathcal{P}),\mu,[\![p_i]\!]\rangle_{p\in\mathcal{P}}$, with and $[\![p]\!]$ defined by: $i[\![p]\!] = i\cap\{j\in\wp(\mathcal{P}) : p\in j\}$. It is almost immediate that $[\![p]\!]$ is introspective, monotone, and stable. This means that concrete models are a special case of the models given by definition 3.2. By abuse of notation we denote these models by $\mathcal{P}$, and allow $\mathcal{P}$ to be infinite.

Concrete models have the advantage of turning an $\mathcal{L}_0$-semantics into one for $\mathcal{L}_1$-sentences. Since a world in an information state is equivalent to an valuation $m : \mathcal{P} \longrightarrow \{0,1\}$, the concrete models are built by taking the power of the set of models for CPL (Classical Propositional Logic). Given the $\mathcal{L}_0$-models, a concrete model contains all possible information states which can be obtained from them. By contrast, definition 3.2 allows models of this kind to consists of a field over a *subset* of the set of all $\mathcal{L}_0$-models.[5] Proposition 3.6 has some properties of updating concrete information states with a sequence Γ of $\mathcal{L}_0$-formulas in terms of their models.

Proposition 3.6 Let i be a state in $\mathcal{P}$, Γ a sequence of $\mathcal{L}_0$-formulas. Set $m(\Gamma) = m(\bigwedge\Gamma)$, m a valuation.

i) $\top[\Gamma] = \{m\in\mathcal{P}\longrightarrow\{0,1\} : m(\Gamma)=1\}$

ii) $i[\Gamma] = \{m\in i : m(\Gamma)=1\}$

iii) $i[\Gamma] = i$ iff $i\subseteq\top[\Gamma]$ iff for all $m\in i : m(\Gamma)=1$

iv) $\top[\Gamma] = \bot$ iff $\Gamma\vdash_{\mathrm{cpl}}$

Proof: by 3.3, and the completeness of classical propositional logic. □

In case $i[\varphi] = i$, φ is *accepted* in i. Proposition 3.6 (iii) shows that acceptance generalizes truth: i accepts φ iff φ is valid in i.

4 Completeness and decidability

For convenience of the reader we repeat the definitions of $\models_{\mathrm{tc}}$, $\models_{\mathrm{ic}}$ and $\models_{\mathrm{utc}}$, but now in the context of $\mathbf{M}$.

[5] A field is a non-empty set of sets which is closed under intersection and complementation.

Definition 4.1 (logical consequence) Let $\mathbf{I}$ a class of models, $\mathcal{I} \in \mathbf{I}$, and let Π, τ be a sequence of $\mathcal{L}_1$-formulas.

- $\Pi \models^{\mathcal{I}}_{\mathrm{tc}} \tau$ iff: for all $i \in \mathcal{I}$: $i[\tau] = i$, if $i[\sigma] = i$ for each σ in Π.
- $\Pi \models^{\mathcal{I}}_{\mathrm{ic}} \tau$ iff: $\top^{\mathcal{I}}[\Pi, \tau] = \top^{\mathcal{I}}[\Pi]$.
- $\Pi \models^{\mathcal{I}}_{\mathrm{utc}} \tau$ iff: $i[\Pi, \tau] = i[\Pi]$ for each $i \in \mathcal{I}$.

For each of the relations $\models$, τ is called a *consequence of* Π relative to $\mathbf{I}$ —$\Pi \models \tau$— iff $\Pi \models^{\mathcal{I}} \tau$ for each $\mathcal{I} \in \mathbf{I}$. □

The updates in $\mathbf{M}$ are Idempotent, but Permutation fails. So, proposition 2.3 does not apply. The next examples show that $\models_{\mathrm{tc}}$, $\models_{\mathrm{ic}}$ and $\models_{\mathrm{utc}}$ are really different:

$$\begin{array}{lll} Mp, \neg p \models_{\mathrm{tc}} \bot & Mp, \neg p \not\models_{\mathrm{ic}} \bot & Mp, \neg p \not\models_{\mathrm{utc}} \bot \\ p \not\models_{\mathrm{tc}} Mq & p \models_{\mathrm{ic}} Mq & p \not\models_{\mathrm{utc}} Mq \end{array}$$

By combining the system $\mathbf{M}$ as defined in section 3 with either of the consequence relations $\models_{\mathrm{i}}$ ($i \in \{\mathrm{tc},\ \mathrm{ic},\ \mathrm{utc}\}$), three logics $\mathbf{M}_{\mathrm{tc}}$, $\mathbf{M}_{\mathrm{ic}}$, and $\mathbf{M}_{\mathrm{utc}}$ are generated. In this section we will take a closer look at these logics, and introduce a complete sequent system for each.

The sequent systems combine a sequent system for the "object" language $\mathcal{L}_0$ with one for the "meta"-language $\mathcal{L}_1$. More in particular, the three might-logics share the system M which consists of two general structural rules (reflexivity, cautious cut) together with classical logical rules for the constants $\neg$ and $\wedge$, and the structural rules monotonicity, contraction, and permutation for $\mathcal{L}_0$-sequents. The system M is extended to a system M_{i} (i as above) by adding logical rules for the 'might'-operator and structural rules for $\mathcal{L}_1$-sequents.

With a view to natural language semantics it seems less than ideal to distinguish between levels of language. But logically it is proficient. For example, the completeness and decidability results below directly extend well-know facts concerning classical logic.

Conventions The letters φ, ψ, $\chi, \ldots$ vary over $\mathcal{L}_0$-formulas, and Δ, $\Gamma, \ldots$ over finite, possibly empty sequences of $\mathcal{L}_0$-formulas. $\mathcal{L}_1$-formulas are denoted by σ, τ, $\mu, \ldots$, and finite, possibly empty sequences of such formulas by Π, $\Lambda, \ldots$. The letters may carry sub- or superscripts. The set $\textsc{prop}(\Pi)$ consists of the proposition letters used to built Π. Π_0 refers to the sequence of $\mathcal{L}_0$-formulas, which results from erasing the M-formulas in Π. A sequent is a pair $\langle \Pi, \sigma \rangle$. The sequent $\langle \Pi, \tau \rangle$ is derivable within sequent system S iff $\Pi \vdash_{\mathrm{S}} \tau$ can be derived from instances of the axioms and the rules of S.

4.1 The common part M

All three logics have the following rules in common.

The Classical Part consists of $\mathcal{L}_0$-sequents.

Logical rules

$$\frac{\Gamma, \varphi_{\mathrm{i}} \vdash \chi}{\Gamma, \varphi_1 \wedge \varphi_2 \vdash \chi}\ \mathrm{L}_{\wedge}^{\mathrm{i}} \qquad \frac{\Gamma \vdash \varphi_1 \quad \Gamma \vdash \varphi_2}{\Gamma \vdash \varphi_1 \wedge \varphi_2}\ \mathrm{R}_{\wedge}$$

$$\frac{\Gamma \vdash \varphi}{\Gamma, \neg\varphi \vdash}\ \mathrm{L}_{\neg} \qquad \frac{\Gamma, \varphi \vdash}{\Gamma \vdash \neg\varphi}\ \mathrm{R}_{\neg} \qquad \frac{\Gamma, \varphi \vdash \chi}{\Gamma, \neg\neg\varphi \vdash \chi}\ \mathrm{L}_{\neg\neg}$$

Structural rules

$$\frac{\Gamma \vdash}{\Gamma \vdash \varphi}\ \mathrm{R}_{\mathrm{mon}} \quad \frac{\Gamma, \Delta \vdash \chi}{\Gamma, \varphi, \Delta \vdash \chi}\ \mathrm{mon} \quad \frac{\Gamma, \varphi, \varphi, \Delta \vdash \chi}{\Gamma, \varphi, \Delta \vdash \chi}\ \mathrm{contr} \quad \frac{\Gamma, \varphi, \psi, \Delta \vdash \chi}{\Gamma, \psi, \varphi, \Delta \vdash \chi}\ \mathrm{perm}$$

The Common $\mathcal{L}_1$ Part consist of $\mathcal{L}_1$-sequents.

$$\frac{}{\sigma \vdash \sigma}\ \mathrm{refl} \qquad \frac{\Pi \vdash \sigma \quad \Pi, \sigma, \Lambda \vdash \tau}{\Pi, \Lambda \vdash \tau}\ \text{cautious cut}$$

Observe that in the context of the classical structural rules cautious cut, referred to as 'ccut', is equivalent to the familiar cut rule.

$$\frac{\Gamma \vdash \varphi \quad \Delta, \varphi, \Delta' \vdash \psi}{\Delta, \Gamma, \Delta' \vdash \psi}\ \mathrm{cut}$$

To be precise, monotonicity and ccut imply cut, while contraction, permutation and cut imply ccut. This means that the system M_0, which is M restricted to $\mathcal{L}_0$-sequents, is a sound and complete sequent calculus for classical logic.

Fact 4.2 $\Gamma \vdash_{M_0} \varphi \Leftrightarrow \Gamma \models_{\mathrm{cpl}} \varphi$

Since the might-logics are obtained by adding logical and structural rules to M, they also have:

Fact 4.3 $\Gamma \vdash_{\mathrm{cpl}} \varphi \Rightarrow \Gamma \vdash_M \varphi$

Finally, for each extension M_{i} of M the reflexivity rule can be restricted to proposition letters. This is handy, for it means that in the ccut-free variant of M_{i} the rule need not be considered as 'might' introducing.

4.2 The test logic $\mathbf{M}_{tc}$

The logic $\mathbf{M}_{\mathrm{tc}}$ is the simplest of the three logics introduced above. In this section we define a sequent system M_{tc} that completely axiomatizes $\mathbf{M}_{\mathrm{tc}}$. M_{tc} consists of M together with the following rules:[6]

Logical rules

$$\frac{\Delta, \varphi \vdash \psi}{\Delta, \mathrm{M}\varphi \vdash \mathrm{M}\psi}\ \mathrm{M} \qquad\qquad \frac{\Pi \vdash \varphi}{\Pi \vdash \mathrm{M}\varphi}\ \mathrm{R_m}$$

Structural rules

$$\frac{\Pi', \Pi'' \vdash \tau}{\Pi', \sigma, \Pi'' \vdash \tau}\ \mathrm{L_{mon}} \qquad \frac{\Pi \vdash}{\Pi \vdash \tau}\ \mathrm{R_{mon}} \qquad \frac{\Pi', \sigma, \rho, \Pi' \vdash \tau}{\Pi', \rho, \sigma, \Pi' \vdash \tau}\ \mathrm{perm}$$

In the soundness proof below, $i[\Pi] = i$ will mean that $i[\sigma] = i$ for all σ in Π. (Given that all updates are introspective, we are allowed to so).

Proposition 4.4 (soundness) The system M_{tc} is sound with respect to class of all models: if $\Pi \vdash_{\mathrm{tc}} \tau$ then $\Pi \models_{\mathrm{tc}} \tau$.
Proof: We show that rule M preserves $\models^{\mathcal{I}}_{\mathrm{tc}}$. The other cases are similar or simpler. Assume that $\Delta, \varphi \models^{\mathcal{I}}_{\mathrm{tc}} \psi$. That is: if $i[\Delta, \varphi] = i$ then $i[\psi] = i$, for each $i \in \mathcal{I}$. We derive a contradiction from the assumption that (i) $j[\Delta] = j$ (ii) $j[\mathrm{M}\varphi] = j$ and (iii) $j[\mathrm{M}\psi] \neq j$, for some $j \in \mathcal{I}$. It follows from (iii) that $j \neq \bot$ and $j[\psi] \neq \bot$. Since $j \neq \bot$ (ii) gives $j[\varphi] \neq \bot$. Set $k = j[\varphi]$. Then $k[\Delta] = k$, due to (i) and the fact that Δ is $\mathcal{L}_0$. Also, $k[\varphi] = k$ by the idempotency of $[\varphi]$. The main assumption gives: $k[\psi] = k$. So, $j[\varphi][\psi] = j[\psi][\varphi] \neq \bot$, which is in conflict with $j[\psi] = \bot$. $\square$

In rule M, the restriction to $\mathcal{L}_0$-formulas is crucial, otherwise soundness would be lost. For let $\mathcal{P} = \{p, q\}$. Then $p \vee q, \mathrm{M}\neg p, \neg q \models^{\mathcal{P}}_{\mathrm{tc}} \bot$. But $p \vee q, \mathrm{M}\neg p, \mathrm{M}\neg q \not\models^{\mathcal{P}}_{\mathrm{tc}} \mathrm{M}\bot$. E.g., $j = \{m_1, m_2\}$ with $m_1(p) = m_2(q) = 1$ and $m_2(p) = m_1(q) = 0$ accepts the last three formulas.

In the next proof $\Pi^\bullet$ stands for the sequence Π with all M-operators erased.

Theorem 4.5 (completeness) If $\Pi \models_{\mathrm{tc}} \tau$ then $\Pi \vdash_{\mathrm{tc}} \tau$.
Proof: Assume $\Pi \not\vdash_{\mathrm{tc}} \tau$. Due to $\mathrm{L_{mon}}$ and $\mathrm{R_m}$ this means that $\Pi_0 \not\vdash_{\mathrm{tc}} \tau^\bullet$. Let $\mathcal{P} \supseteq \textsc{prop}(\Pi, \tau)$. Classical completeness gives a model $m : \mathcal{P} \longrightarrow \{0, 1\}$ with (*) $m(\Pi_0) = 1$ and $m(\Pi_0, \tau^\bullet) = 0$. We use this to find an $i \in \mathcal{P}$ with $i[\Pi] = i$ and $i[\tau] \neq i$.

[6] We use the same name for possibly different rules in different logics; it is the obvious choice and no confusion is likely.

• $\tau \in \mathcal{L}_0$. Consider $i = \top^{\mathcal{P}}[\Pi_0]$. Since $i \ni m \notin i[\tau]$ by (*) and proposition 3.6 (iii), it holds that $i[\tau] \neq i$,. But also $i[\sigma] = i$ for the formulas σ in Π. In case σ is in Π_0 this is clear from proposition 3.6 (iii). So let $\Pi \equiv \Pi', \mathrm{M}\varphi, \Pi''$. It is sufficient to prove $i[\varphi] \neq \bot$. If not proposition 3.6 (iv) gives: $\Pi_0, \varphi \vdash_{\mathrm{cpl}}$. Rule M with ψ (and hence $\mathrm{M}\psi$) empty: $\Pi_0, \mathrm{M}\varphi \vdash_{\mathrm{tc}}$. So by permutation: $\Pi_0', \mathrm{M}\varphi, \Pi_0'' \vdash_{\mathrm{tc}}$. $\mathrm{L_{mon}}$ and $\mathrm{R_{mon}}$ turn this into the contradictory: $\Pi \vdash_{\mathrm{tc}} \tau$.

• $\tau \notin \mathcal{L}_0$. Consider $i = \top^{\mathcal{P}}[\neg\tau^\bullet, \Pi_0]$. By proposition 3.6 (iii): $i[\tau^\bullet] = \bot$. So $i[\tau] = \bot \neq i$, using (*). But $i[\sigma] = i$ for σ in Π. As before the σ in Π_0 pose no problem. In case $\Pi \equiv \Pi', \mathrm{M}\varphi, \Pi'$ it suffices to prove $i[\varphi] \neq \bot$. Assume otherwise. Proposition 3.6 (iv) gives: $\neg\tau^\bullet, \Pi_0, \varphi \vdash_{\mathrm{cpl}}$. Hence $\Pi_0, \varphi \vdash_{\mathrm{cpl}} \tau^\bullet$. Contrary to the assumption one can show that $\Pi \vdash_{\mathrm{tc}} \tau$ by means of M, permutation, and the monotonicity rules. □

The next two corollaries are seen to hold by checking the above proofs.

Corollary 4.6 M_{tc} is sound and complete with respect to the model $\mathcal{P}$, with $\mathcal{P}$ the proposition letters used to generate the formulas. □

Corollary 4.7 (ccut elimination) The system M_{tc} is deductively equivalent to the system without the ccut rule. □

Corollary 4.8 (decidability) The logic M_{tc} is decidable.
Proof: For the ccut-free version of M_{tc} it holds that $\Pi \vdash_{\mathrm{tc}} \tau$ iff:

(i) $\Pi^{\mathrm{p}} \vdash_{\mathrm{cpl}}$ or $\Pi^{\mathrm{p}} \vdash_{\mathrm{cpl}} \tau^\bullet$ for a permutation Π^{p} of Π; or

(ii) $(\Pi^{\mathrm{p}})_0 \vdash_{\mathrm{cpl}} \tau^\bullet$ for a permutation Π^{p} of Π.

But (i) and (ii) only use finitely many instances of the decidable relation $\vdash_{\mathrm{cpl}}$. □

4.3 Starting from ignorance: the logic $\mathbf{M}_{ic}$

The main result of this section states for each finite set of atoms $\mathcal{P}$ we can find a sequent system M_{ic} that completely axiomatizes $\mathbf{M}_{\mathrm{ic}}$ on the model $\mathcal{P}$. The sequent system M_{ic} over the language $\mathcal{P}$ consists of M and the following $\mathcal{L}_1$-rules.

Logical rules

$$\frac{\Pi, \varphi \vdash}{\Pi, \mathrm{M}\varphi \vdash}\,\mathrm{L_{m1}} \qquad \frac{\Pi, \Lambda \vdash \tau}{\Pi, \mathrm{M}\varphi, \Lambda \vdash \tau}\,\mathrm{L_{m2}} \qquad \frac{\Pi \vdash \varphi}{\Pi \vdash \mathrm{M}\varphi}\,\mathrm{R_m}$$

$$\frac{\delta \vdash \varphi_1, \ldots, \delta \vdash \varphi_n, \delta \vdash \psi}{\varphi_1, \ldots, \varphi_n \vdash \mathrm{M}\psi}\ \mathrm{M}$$

In rule M, δ is required to be a complete diagram relative to the finite $\mathcal{P}$; that is, δ is a sequence $(\neg)p_1, \ldots, (\neg)p_n$ of (negations of) atoms such that each $p \in \mathcal{P}$ occurs exactly once.

Structural rules

$$\frac{\Pi \vdash}{\Pi, \Lambda \vdash}\ \mathrm{L_{mon}} \qquad\qquad \frac{\Pi \vdash}{\Pi \vdash \sigma}\ \mathrm{R_{mon}}$$

Proposition 4.9 (soundness) The system M_{ic} is sound with respect to the model $\mathcal{P}$: if $\Pi \vdash_{\mathrm{ic}} \tau$ then $\Pi \models^{\mathcal{P}}_{\mathrm{ic}} \tau$.
Proof: We only consider rule M. Assume $\top[\delta, \gamma] = \top[\delta]$ for all γ in Γ, ψ. Since δ is a diagram, classical completeness gives a valuation m such that $m(\Gamma, \psi) = 1$. So $\top[\Gamma, \psi] \neq \bot$, and therefore $\top[\Gamma, \mathrm{M}\psi] = \top[\Gamma]$. □

Notice that the proof requires $\top$ to contain all $\mathcal{L}_0$-models. Completeness is established by means of the following lemma.

Lemma 4.10 If $\Pi \not\vdash_{\mathrm{ic}}$, then $\top[\Pi'] = \top[\Pi'_0]$ for each initial Π' of Π.
Proof: We distinguish two cases in an induction on the length of Π'.
• $\Pi' \equiv \Lambda, \varphi$ with $\varphi \in \mathcal{L}_0$. Then $\top[\Lambda, \varphi] =_{\mathrm{i.h.}} \top[\Lambda_0, \varphi] = \top[(\Lambda, \varphi)_0]$.
• $\Pi' \equiv \Lambda, \mathrm{M}\varphi$. It is sufficient to prove $\top[\Lambda, \varphi] \neq \bot$. For then $\top[\Lambda, \mathrm{M}\varphi] = \top[\Lambda] = \top[(\Lambda, \mathrm{M}\varphi)_0]$. So let $\Pi \equiv \Pi', \Pi''$, and assume $\top[\Lambda, \varphi] = \bot$. By the induction hypothesis: $\top[\Lambda_0, \varphi] = \bot$. Proposition 3.6 (iv): $\Lambda_0, \varphi \vdash_{\mathrm{cpl}}$. So $\Pi', \Pi'' \vdash_{\mathrm{ic}}$ by respectively using $\mathrm{L_{m1}}$, $\mathrm{L_{mon}}$, and $\mathrm{L_{m2}}$. But this contradicts the assumption. □

Theorem 4.11 (completeness) If $\Pi \models^{\mathcal{P}}_{\mathrm{ic}} \tau$, then $\Pi \vdash_{\mathrm{ic}} \tau$.
Proof: Assume $\Pi \not\vdash_{\mathrm{ic}} \tau$. Two cases are to be distinguished. • $\tau \in \mathcal{L}_0$. Then with $\mathrm{L_{m2}}$: $\Pi_0 \not\vdash_{\mathrm{ic}} \tau$, and hence $\Pi_0 \not\vdash_{\mathrm{cpl}} \tau$. Classical completeness gives some m with $m(\Pi_0) = 1$ and $m(\Pi_0, \tau) = 0$. Due to $\mathrm{R_{mon}}$ $\Pi \not\vdash_{\mathrm{ic}}$. Hence, by lemma 4.10: $\top[\Pi] = \top[\Pi_0] \ni m \notin \top[\Pi_0, \tau] = \top[\Pi, \tau]$
• $\tau \equiv \mathrm{M}\psi$. Claim: $\Pi_0, \psi \vdash_{\mathrm{cpl}}$. To prove this claim, note that if $\Pi_0, \psi \not\vdash_{\mathrm{cpl}}$ classical completeness gives a diagram δ with $\delta \vdash_{\mathrm{ic}} \gamma$ for each γ in Π_0, ψ. Rule M: $\Pi_0 \vdash_{\mathrm{tp}} \mathrm{M}\psi$. Hence, by $\mathrm{L_{m2}}$: $\Pi \vdash_{\mathrm{tp}} \mathrm{M}\psi$, a contradiction.

Since $\Pi_0, \psi \vdash_{\mathrm{cpl}}$ and by $\mathrm{R_{mon}}$: $\Pi \not\vdash_{\mathrm{ic}}$, lemma 4.10 and proposition 3.6 (iv) imply: $\top[\Pi, \psi] = \top[\Pi_0, \psi] = \bot$. Hence: $\top[\Pi, \mathrm{M}\psi] = \bot$. On the other hand, by lemma 4.10 and classical completeness: $\top[\Pi] = \top[\Pi_0] \neq \bot = \top[\Pi, \mathrm{M}\psi]$.

In both cases it is found that $\Pi \not\models^{\mathcal{P}}_{\mathrm{ic}} \tau$. □

Since $\mathcal{P}$ is finite, it is clear that M_{ic} is decidable.

4.4 The update-test logic $\mathbf{M}_{utc}$

In this section we axiomatize $\mathbf{M}_{\rm utc}$, i.e., $\mathbf{M}$ with $\models_{\rm utc}$ as its consequence relation. The sequent system $M_{\rm utc}$ extends M with two logical rules:

$$\frac{\Pi, \Lambda \vdash \tau}{\Pi, \mathrm{M}\varphi, \Lambda \vdash \tau}\ \mathrm{L_m} \qquad\qquad \frac{\Delta, \varphi \vdash \chi\ \textit{all}\ \chi\ \textit{in}\ \Gamma, \psi)}{\Delta, \mathrm{M}\varphi, \Gamma \vdash \mathrm{M}\psi}\ \mathrm{M}$$

Rule M allows φ, Γ, and ψ to be empty. Therefore, the following rules in are instances of M.

$$\frac{\Delta, \varphi \vdash \psi}{\Delta, \mathrm{M}\varphi \vdash \mathrm{M}\psi} \qquad\qquad \frac{\Delta \vdash \chi\ (\textit{all}\ \chi\ \textit{in}\ \Gamma, \psi)}{\Delta, \Gamma \vdash \mathrm{M}\psi}$$

Proposition 4.12 (soundness) The system $M_{\rm utc}$ is sound with respect to class of all models: if $\Pi \vdash_{\rm utc} \tau$ then $\Pi \models_{\rm utc} \tau$.
Proof: Again we discuss only the rule M. Assume $\Delta, \varphi \models^{\mathcal{I}}_{\rm utc} \gamma$ for each γ in Γ, ψ. Pick $i \in \mathcal{I}$. In case $i[\Delta, \varphi] = \perp$ it is clear that $i[\Delta, \mathrm{M}\varphi, \Gamma, \mathrm{M}\psi] = i[\Delta, \mathrm{M}\varphi, \Gamma]$. So let $i[\Delta, \varphi] \neq \perp$. By assumption $i[\Delta, \varphi] = i[\Delta, \varphi, \Gamma, \psi]$. So $i[\Delta, \varphi, \Gamma, \psi] \neq \perp$ and hence $i[\Delta, \Gamma, \psi] \neq \perp$. Therefore (*) $i[\Delta, \Gamma, \mathrm{M}\psi] = i[\Delta, \Gamma]$. And since $i[\Delta, \varphi] \neq \perp$, also (**) $i[\Delta, \mathrm{M}\varphi] = i[\Delta]$. But (*) and (**) imply $i[\Delta, \mathrm{M}\varphi, \Gamma, \mathrm{M}\psi] = i[\Delta, \mathrm{M}\varphi, \Gamma]$. The choice of i was arbitrary, so $\Delta, \mathrm{M}\varphi, \Gamma \models^{\mathcal{I}}_{\rm utc} \mathrm{M}\psi$. $\Box$

Lemma 4.13 (consistency lemma) If $\Pi \not\vdash_{\rm utc} \tau$, then for each initial segment Π' of Π and each $\mathcal{P} \supseteq \textsc{prop}(\Pi, \tau)$: $\top^{\mathcal{P}}[\neg \bigwedge(\Pi_0, \tau^\bullet), \Pi'] \neq \perp$.
Proof: Let $\mathcal{P}$ be a model of the relevant kind and set $i = \top^{\mathcal{P}}[\neg \bigwedge(\Pi_0, \tau^\bullet)]$. We use induction on the length of $\Pi' \equiv \Lambda, \sigma$.

Observe that it is sufficient to prove $i[\Lambda, \sigma^\bullet] \neq \perp$. For if $\sigma^\bullet \equiv \sigma$ we are done. Whereas if $\sigma \equiv \mathrm{M}\varphi$ we have: $i[\Pi'] = i[\Lambda, \mathrm{M}\varphi] = i[\Lambda] \neq_{\rm i.h.} \perp$. So assume $i[\Lambda, \sigma^\bullet] = \perp$. By the i.h. and corollary 3.5 (v): $i[\Lambda_0, \sigma^\bullet] = \perp$. According to corollary 3.6 and the definition of i: $\neg\bigwedge(\Pi_0, \tau^\bullet), \Lambda_0, \sigma^\bullet \vdash_{\rm cpl}$. So by classical reasoning and fact 4.3: $\Lambda_0, \sigma^\bullet \vdash_{\rm utc} \bigwedge(\Pi_0, \tau^\bullet)$. Let $\Pi \equiv \Pi', \Pi''$. $\mathrm{R}_\wedge$ yields: $\Lambda_0, \sigma^\bullet \vdash_{\rm utc} \gamma$ for each γ in $\Pi''_0, \tau^\bullet$. Rule M: $\Lambda_0, \sigma, \Pi''_0 \vdash_{\rm utc} \tau$. (To be precise, if $\sigma \in \mathcal{L}_0$ we assume φ in M to be empty, and similarly for τ.) $\mathrm{L_m}$ proves: $\Pi \vdash_{\rm utc} \tau$, a contradiction. $\Box$

Theorem 4.14 (completeness) If $\Pi \models_{\rm utc} \tau$ then $\Pi \vdash_{\rm utc} \tau$.
Proof: Assume $\Pi \not\vdash_{\rm utc} \tau$, and let $\mathcal{P} \supseteq \textsc{prop}(\Pi, \tau)$. Set $i = \top^{\mathcal{P}}[\neg \bigwedge(\Pi_0, \tau^\bullet)]$. Then by lemma 4.13 and corollary 3.5 (v): $i[\Pi] = i[\Pi_0] = \top[\neg\tau^\bullet, \Pi_0] \neq \perp$. It follows that $i[\Pi, \tau^\bullet] = \top[\neg\tau^\bullet, \Pi_0, \tau^\bullet] = \perp$. Whether or not $\tau^\bullet = \tau$, in both cases we get $i[\Pi] \neq \perp = i[\Pi, \tau]$. Therefore: $\Pi \not\models^{\mathcal{P}}_{\rm utc} \tau$. $\Box$

A check of the above proofs gives some corollaries.

Corollary 4.15 M_{utc} is sound and complete with respect to the model $\mathcal{P}$, with $\mathcal{P}$ the proposition letters used to generate the formulas. $\square$

Corollary 4.16 If $\mathcal{P} = \textsc{prop}(\Pi, \tau)$, then $\Pi \vdash_{\mathrm{utc}} \tau$ iff $\Pi \models^{\mathcal{P}}_{\mathrm{utc}} \tau$. $\square$

Corollary 4.17 (decidability) The logic $\mathbf{M}_{\mathrm{utc}}$ is decidable.
Proof: In order to check whether or not $\Pi \vdash_{\mathrm{utc}} \tau$ it suffices to search the finitely many states of $\textsc{prop}(\Pi, \tau)$ (corollary 4.16). As soon as a counter model is found we know $\Pi \not\vdash_{\mathrm{utc}} \tau$, but otherwise: $\Pi \vdash_{\mathrm{utc}} \tau$. $\square$

Corollary 4.18 (ccut elimination) The cautious cut rule can be eliminated from M_{utc}. $\square$
Proof: If at all, cautious cut is only used in the classical part, where it is eliminable. $\square$

This ends our discussion of $\mathbf{M}_{\mathrm{utc}}$.

We have been careful in presenting M as an extension of classical propositional logic. But to what extent does this approach generalize to other 'base' logics? That is, is it possible to restate the above result as a preservation result of the form: for each complete $\mathcal{L}_0$-logic of a certain kind, there exists a might logic which is complete (and similarly for other properties). The main point seems to be to find a generalization of the concrete models. We have to leave this question open. A similar question can be asked with respect to cautious cut elimination. But here we need not bother about semantical issues since the result can be proved syntactically.

5 Cautious Cut Elimination

In this section we forget about set-theoretic interpretations and confine ourselves to syntactic methods. We shall prove the following theorem concerning generalizations of the update test logic for 'might'.

Theorem 5.1 Let $\vdash_0$ be a consequence relation for $\mathcal{L}_0$-sequents which is reflexive, and closed under monotony and cautious cut. Extend the language to an $\mathcal{L}_1$-language as in definition 3.1, and extend $\vdash_0$ to $\vdash_1$ for the $\mathcal{L}_1$-language by closure under cautious cut and the rules M and $\mathrm{L_m}$ of section 4.4. If $\vdash_0$ has cautious cut elimination, then so has $\vdash_1$.

Notice that we need not assume reflexivity for $\mathcal{L}_1$-sequents, since it can be derived by means of rule M. This is handy, for it means that in the ccut-free variant of $\vdash_1$ reflexivity need not be considered as 'might' introducing.

Proof of theorem 5.1 As in case of $\vdash_{\mathrm{m}}$ the relation $\vdash_0$ will contain logical and structural rules for $\mathcal{L}_0$-sequents. But the use of these rules is blocked after an application of M or $\mathrm{L_m}$. This means that if ccut is applied to ccut-free premisses that part of a derivation will have the following structure:

$$\dfrac{\dfrac{\dfrac{\Delta, \varphi \vdash_0 \gamma \ (\gamma \text{ in } \Gamma, \sigma^\bullet)}{\Delta, \mathrm{M}\varphi, \Gamma \vdash \sigma}\,(\mathrm{M})}{\Pi \vdash \sigma}\,(\mathrm{L_m})^* \qquad \dfrac{\dfrac{\Delta', \psi \vdash_0 \chi \ (\chi \text{ in } \Gamma', \tau^\bullet)}{\Delta', \mathrm{M}\psi, \Gamma' \vdash \tau}\,(\mathrm{M})}{\Pi, \sigma, \Lambda \vdash \tau}\,(\mathrm{L_m})^*}{\Pi, \Lambda \vdash \tau}\,\text{ccut}$$

Here, (M) indicates that M is applied at most once, and $(\mathrm{L_m})^*$ that $\mathrm{L_m}$ is used finitely many (possibly zero) times. Given this general form we prove ccut-elimination as follows.

Let $\mathcal{D}$ be a derivation for $\vdash_1$. If $\mathcal{D}$ is ccut-free we are done. Otherwise select an occurrence of ccut with cutt-free premisses. If this occurrence lies within the $\mathcal{L}_0$-part of $\mathcal{D}$ we know by assumption how to eliminate it. But if the ccut is applied to $\mathcal{L}_1$-sequents we distinguish four cases.

Case I: There are no applications of M above the cut. Then, the situation is:

$$\dfrac{\dfrac{\Pi_0 \vdash_0 \varphi}{\Pi \vdash \varphi}\,(\mathrm{L_m})^* \qquad \dfrac{\Pi_0, \varphi, \Lambda_0 \vdash_0 \psi}{\Pi, \varphi, \Lambda \vdash \psi}\,(\mathrm{L_m})^*}{\Pi, \Lambda \vdash \psi}\,\text{ccut}$$

This can be reduced to:

$$\dfrac{\dfrac{\Pi_0 \vdash_0 \varphi \quad \Pi_0, \varphi, \Lambda_0 \vdash_0 \psi}{\Pi_0, \Lambda_0 \vdash \psi}\,\text{ccut}}{\Pi, \Lambda \vdash \psi}\,(\mathrm{L_m})^*$$

Here ccut occurs in the $\mathcal{L}_0$-part and is hence eliminable.

Case II: In deriving the right-hand side premiss of the ccut, M is applied once. The situation is:

$$\dfrac{\dfrac{\Pi_0 \vdash_0 \chi}{\Pi \vdash \chi}\,(\mathrm{L_m})^* \qquad \dfrac{\dfrac{\Pi_0' \vdash_0 \gamma \ (\gamma \text{ in } \Lambda_0', \tau^\bullet)}{\Pi_0', \mathrm{M}\varphi, \Lambda_0' \vdash \tau}\,\mathrm{M}}{\Pi, \chi, \Lambda \vdash \tau}\,(\mathrm{L_m})^*}{\Pi, \Lambda \vdash \tau}\,\text{ccut}$$

We distinguish two subcases. When χ occurs in Λ_0' we obtain a ccut-free derivation from the left premiss by deleting this occurrence. But when χ occurs in Π_0' we have $\Pi_0' \equiv \Pi_0, \chi, \Pi_0''$ for some Π_0'' (and hence $\Lambda_0 \equiv \Pi_0'', \Lambda_0'$). Then the above can be reduced to:

$$\cfrac{\cfrac{\cfrac{\Pi_0 \vdash_0 \chi \quad \Pi_0, \chi, \Pi_0'' \vdash_0 \gamma \ (\gamma \text{ in } \Lambda_0', \tau^\bullet)}{\Pi_0, \Pi_0'', \varphi \vdash \gamma \ (\gamma \text{ in } \Lambda_0', \tau^\bullet)}\ \text{ccut}}{\Pi_0, \Pi_0'', \mathrm{M}\varphi, \Lambda_0' \vdash \tau}\ \mathrm{M}}{\Pi, \Lambda \vdash \tau}\ (\mathrm{L_m})^*$$

Again, these ccuts are eliminable by assumption.

Case III: M is used once in deriving the left-hand side premiss of the ccut. This case is trivial, for in the right premiss the ccut formula comes from $\mathrm{L_m}$.

Case IV: The derivations of both premisses contain an application of M. Again the trivial reduction of the previous case may apply, but the situation may also be more interesting:

$$\cfrac{\cfrac{\cfrac{\Pi_0', \varphi \vdash_0 \gamma \ (\gamma \text{ in } \Pi_0'', \psi)}{\Pi_0', \mathrm{M}\varphi, \Pi_0'' \vdash \mathrm{M}\psi}\ (\mathrm{M})}{\Pi', \mathrm{M}\varphi, \Pi'' \vdash \mathrm{M}\psi}\ (\mathrm{L_m})^* \quad \cfrac{\cfrac{\Pi_0', \Pi_0'', \psi \vdash_0 \chi \ (\chi \text{ in } \Lambda_0, \tau^\bullet)}{\Pi_0', \Pi_0'', \mathrm{M}\psi, \Lambda_0 \vdash \tau}\ (\mathrm{M})}{\Pi', \mathrm{M}\varphi, \Pi'', \mathrm{M}\psi, \Lambda \vdash \tau}\ (\mathrm{L_m})^*}{\Pi', \mathrm{M}\varphi, \Pi'', \Lambda \vdash \tau}\ \text{ccut}$$

This reduces to:

$$\cfrac{\cfrac{\cfrac{\Pi_0', \varphi \vdash_0 \chi \ (\chi \text{ in } \Pi_0'', \psi) \quad \cfrac{\Pi_0', \Pi_0'', \psi \vdash_0 \chi \ (\chi \text{ in } \Lambda_0, \tau^\bullet)}{\Pi_0', \varphi, \Pi_0'', \psi \vdash_0 \chi \ (\chi \text{ in } \Lambda_0, \tau^\bullet)}\ \text{mon}}{\Pi_0', \varphi \vdash_0 \chi \ (\chi \text{ in } \Pi_0'', \Lambda_0, \tau^\bullet)}\ \text{ccut}}{\Pi_0', \mathrm{M}\varphi, \Pi_0'', \Lambda \vdash \tau}\ \mathrm{M}}{\Pi', \mathrm{M}\varphi, \Pi'', \Lambda \vdash \tau}\ (\mathrm{L_m})^*$$

This completes the proof of theorem 5.1. □

Note that in reducing the ccut to $\mathcal{L}_0$-sequents the $\mathcal{L}_1$ part of the proof grows at most $n+1$ steps, where n is the length of Π_0'', ψ in the reduction of Case *IV*. The other reductions shorten or do not alter the length of the proof.

If $\vdash_0$ in theorem 5.1 is decidable, $\vdash_1$ can be shown to be decidable too. Except for ccut, the rules of $\vdash_1$ satisfy the subformula property. So the following algorithm to check whether or not $\Pi \vdash_1 \tau$ is recursive:

i) If $\tau \in \mathcal{L}_0$ check whether $\Pi_0 \vdash_0 \tau$.

ii) If $\tau \equiv \mathrm{M}\psi$ check whether $\Pi_0', \varphi \vdash_0 \gamma$ for each γ in Π_0'', ψ, and each partition $\Pi \equiv \Pi', \mathrm{M}\varphi, \Pi''$ (φ, Π_0'', or ψ may be empty).

By assumption $\vdash_0$ is decidable, so the recipe defines a finite search space with all possible initial sequents to introduce the M-formulas in Π, τ. Therefore, $\Pi \vdash_1 \tau$ iff the algorithm finds a derivable $\mathcal{L}_0$-sequent from which $\Pi \vdash_1 \tau$ can be derived. In particular, since $\vdash_{\mathrm{cpl}}$ is decidable this argument gives a syntactic proof of corollary 4.17.

6 Further Issues

In this section we mention two topics for further study.

Firstly, one would like to obtain similar results for formulas with nested occurrences of the might-operator (cf. [Van Eijck and De Vries 1995]). Such nestings are not allowed here, since the reflexivity axiom would then be lost. E.g., the formula $\mathrm{M}p \wedge \neg p$ is not reflexive. One way to go would be to assume that reflexivity only holds for proposition letters, and to argue that the formulas which do not preserve this property are somehow inadmissible. For instance, the example given corresponds to the unacceptable sentence: 'it might be p and it isn't p.'

Secondly, one may wonder about the minimal algebraic structure for the $\mathcal{L}_0$-part. For instance, do we retain completeness and decidability if we generalize the structures to those of the form $\langle I, \wedge, \bot \rangle$ with $\wedge$ associative and idempotent, and $\bot$ a left and right neutral element? [Kanazawa 1994b] has some results in this direction for a partial version of 'might'.

References

Benthem, J. van: 1991a, 'General Dynamics'. *Theoretical Linguistics*, **17**, 159–201.

Benthem, J. van: 1991b, *Language in Action*. North-Holland, Amsterdam.

Benthem, J. van: 1991c, 'Logic and the Flow of Information' In: D. Prawitz, B. Skyrms, and D. Westerståhl (eds.) *Proceedings of the Ninth International Congress of Logic, Methodology and Philosophy of Science, Uppsala 1991.* North-Holland, Amsterdam.

Does, J.M. van der: 1994, Cut *might* cautiously. Bulletin of the IGPL, **3**:2/3.

Eijck, J. van, and F.J. de Vries: 1995, 'Reasoning about Update Logic' *Journal of Philosophical Logic*, **24**, 19–45.

Groenendijk, J. and M. Stokhof: 1991. Dynamic predicate logic. *Linguistics and Philosophy*, **14**, 39–100.

Groeneveld, W.: 1995, *Logical Investigations into Dynamic Semantics.* Dissertation, ILLC, Department of Philosophy, University of Amsterdam.

Groeneveld, W. and F. Veltman: 1994, 'Inference Systems for Update Semantics' In: R. Cooper and J. Groenendijk (eds.), *Integrating Semantic Theories II.* DYANA report R2.1B, Amsterdam.

Kanazawa, M.: 1994a, 'Completeness and Decidability of the Mixed Style of Inference with Composition.' In : P. Dekker, and M. Stokhof: 1994, *Proceedings of the Ninth Amsterdam Colloquium.* ILLC, University of Amsterdam.

Kanazawa, M.: 1994b, 'Comments on 'Inference Systems for Update Semantics' by Willem Groeneveld and Frank Veltman' In: R. Cooper and J. Groenendijk (eds.), *Integrating Semantic Theories II.* DYANA report R2.1B, Amsterdam.

Veltman, F.: 1991, 'Defaults in Update Semantics' ILLC prepublication series, LP–91–02. To appear in *Journal of Philosophical Logic*.