Epistemic Logic with Relevant Agents

Ondrej Majer Michal Peliš*

1 Introduction

The aim of epistemic logics is to formalize epistemic states and actions of (possibly human) rational agents. A traditional means for representing these states and actions employs the framework of modal logics, where knowledge corresponds to some necessity operator. Modal axioms (K, T, 4, 5,…) then correspond to structural properties of the agent's knowledge. Employing strong modal systems such as S5 leads to representations of agents who are too ideal in many respects — they are logically omniscient, they have a perfect reflection of their both positive and negative knowledge (positive and negative introspection) etc. Sometimes these representations are called epistemic logics of *potential* rather than actual knowledge.

Frameworks representing only perfect agents have been frequently criticized, see (Fagin, Halpern, Moses, & Vardi, 2003) and (Duc, 2001), and some steps towards more realistic representations have been made (e.g., (Duc, 2001)). We also attempt to represent agents in an environment more realistically. Our motivation is epistemic, we shall concentrate on an agent working with experimental scientific data.

A realistic agent

Our agent is a scientist undertaking experiments or observations. Her typical environment is an experimental setup and her knowledge is usually experimental data (inputs and outputs of an experiment/observation) and some generalizations extracted from the experimental data.

* Work on this text was supported in part by grant no. 401/07/0904 of the Grant Agency of the Czech Republic and in part by grant no. IAA900090703 (Dynamic formal systems) of the Grant Agency of the Academy of Sciences of the Czech Republic. We wish to thank to Timothy Childers for valuable comments.

We assume the observations ('facts') are typically represented by atoms and their conjunctions and disjunctions, while generalizations ('regularities') are represented by conditionals (and their combinations). A conditional is supposed to record a regularly observed connection between the facts represented by the antecedent and the facts represented by the consequent.

It seems to be clear that for many reasons the material implication is not an appropriate representation of such a conditional. One of the main reasons is that the material implication may connect any two arbitrary formulas α, β. For example,

1. $\alpha \to (\beta \to \alpha)$,
2. $(\alpha \wedge \neg\alpha) \to \beta$,
3. $\alpha \to (\beta \vee \neg\beta)$,

are tautologies of classical logic. In our epistemic interpretation the material implication would make a 'law' from every two 'facts', which would obviously make the representation useless. It has other undesirable properties. It cannot deal with errors in the data, which result to contradictory facts (a situation which may very well happen in the scientific practice due to equipment errors). One such error corrupts all the remaining data (from a contradiction everything follows — see 2). It also admits 'laws' which are of no use as their consequent is a tautology (as in 3 — a tautology follows from anything)

The tautologies 1–3 are just examples of the paradoxes of material implication. As these 'paradoxes' were completely solved only in the systems of relevant logics, the obvious choice for a conditional for our scientific agent is relevant implication.

2 Relational semantics for relevant logics

Our point of departure will be the distributive relevant logic R of Anderson and Belnap (1975). The most natural way to introduce relevant logics is certainly proof theoretical (see, e.g., (Paoli, 2002)). However we would like to follow the modal tradition in representing an agent's epistemic states as a set of formulas and make the agent's knowledge dependent not only on the current epistemic state, but also on the states epistemic alternatives. Technically speaking we want to use a relational semantics. This cannot be a standard Kripke semantics with possible worlds and a binary accessibility relation, but a more general relational structure.

Formally our framework will be based on the Routley–Meyer semantics, as developed by Mares (Mares, 2004), Restall (Restall, 1999), Paoli (Paoli, 2002), and others, to which we shall add epistemic modalities. This semantics has been under constant attack for its seeming unintuitivness, but we believe it fits very well our motivations.

We give an informal exposition of structures in the relevant frame and definition of connectives (for formal definitions see the appendix A).

Relevant frame

A relevant frame is a structure $\mathbf{F} = \langle S, L, C, \trianglelefteq, R\rangle$, where S is a non-empty set of situations (states), $L \subseteq S$ is a non-empty set of designated *logical situations*, $C \subseteq S^2$ is a *compatibility* relation, $\trianglelefteq \subseteq S^2$ is a relation of *involvement*, $R \subseteq S^3$ is a *relevance* relation.

A model $\mathbf{M}$ is a relevant frame with the relation $\Vdash$, where $s \Vdash \varphi$ has the same meaning as in Kripke frames — that s carries the information that the formula φ is true ($\varphi \in s$ if we consider states to be sets of formulas).

Situations Situations or information states play the same role as possible worlds in Kripke frames. We assume, they consist of data immediately available to the agent. Like possible worlds, we can see situations as sets of formulas, but, unlike possible worlds, situations might be incomplete (neither φ nor $\neg\varphi$ is true in s) or inconsistent (both φ and $\neg\varphi$ are true in s).

Conjunction and disjunction Classical (weak) conjunction and disjunction correspond to the situation when the agent combines data immediately available to her, i. e. data from her current situation. They behave in the same way as in the case of classical Kripke frames — their validity is given locally:

$$s \Vdash \psi \wedge \varphi \text{ iff } s \Vdash \psi \text{ and } s \Vdash \varphi$$
$$s \Vdash \psi \vee \varphi \text{ iff } s \Vdash \psi \text{ or } s \Vdash \varphi$$

Weak connectives are the only ones which are defined locally. The truth of negation and implication depends also on the data in situations, related to the actual ones, so they are modal by nature. It is possible to define strong conjunction and disjunction as well (see appendix A).

Implication Implication is a modal connective in the sense that its truth depends not only on the current situation, but also on its neighborhood. It can be again understood in analogy with the standard modal reading. We say that an implication $(\varphi \rightarrow \psi)$ holds necessarily in a Kripke frame iff in all worlds where the antecedent holds, the consequent holds as well. In other words, the implication $(\varphi \rightarrow \psi)$ holds through all the neighborhood of the actual world. In the relevant case the neighborhood of a situation s is given by pairs of situations y, z such that s, y, z are related by the ternary relation R. We shall call y, z antecedent and consequent situations, respectively. We say that the implication $(\varphi \rightarrow \psi)$ holds at the situation s iff it is the case that for every antecedent situation y where φ (the antecedent of the implication) holds, ψ (the consequent of the implication) holds at the corresponding consequent situation z.

$$s \Vdash (\varphi \rightarrow \psi) \quad \text{iff} \quad (\forall y, z)(Rsyz \text{ implies } (y \Vdash \varphi \text{ implies } z \Vdash \psi))$$

The relation R reflects in our interpretation actual experimental setups. Antecedent situations correspond to some initial data (outcome of measurements or observations) of some experiment, while the related consequent situations correspond to the corresponding resulting data of the experiment. Implication then corresponds to some (simple) kind of a rule: if I observe in my current situation, that at every experiment (represented by a couple antecedent–consequent situation) each observation of φ is followed by an observation of ψ, then I accept 'ψ follows φ' as a rule.

Logical situations The framework we presented so far is very weak: there are just few tautologies valid in all situations and some of the important ones — those being usually considered as basic logical laws — are missing. For example the widely accepted identity axiom $(\alpha \rightarrow \alpha)$ and the Modus Ponens rule fail to hold in every situation.

This is connected to the question of truth in a relevance frame (model). If we take a hint from Kripke frames, we should equate truth in a frame with truth in every situation. But this would gives us an extremely weak system with some very unpleasant properties (cf. (Restall, 1999)). Designers of relevant logics took a different route — instead of requiring truth in all situations, they identify the truth in a frame just with the truth in all logically well behaved situations. These situations are called *logical.* In order to satisfy the 'good behavior' of a situation l it is enough to require that all the information in any antecedent situation related to l is contained in the corresponding consequent situation as well: for each $x, y \in S$, $Rlxy$ implies $|x| \subseteq |y|$, where $|s|$ is the set of all formulas, which are true in the situation s.

It is easy to see that situations constrained in this way validate both the identity axiom and (implicative) Modus Ponens.

Involvement Involvement is a relation resembling the persistence relation in intuitionistic logic — we can see it as a relation of information growth. However not every two situations which are in inclusion with respect to the validated formulas are in the involvement relation. We require that such an inclusion is observed or witnessed. Not every situation can play the role of the witness — only the logical situations can.

$$x \trianglelefteq y \quad \text{iff} \quad (\exists l \in L)(Rlxy)$$

Negation In Kripke models the *negation* of a formula φ is true at a world iff φ is not true there. As situations can be incomplete and/or inconsistent, this is not an option any more. Negation becomes a modal connective and its meaning depends on the worlds related to the given world by a binary modal relation C known as *compatibility*. Informally we can see the compatible situations as information sources our scientist wants to be consistent with. (Imagine the data of research groups working on related subjects.)

The formula $\neg\varphi$ holds at $s \in S$ iff it is not 'possible' (in the standard modal sense with respect to the relation C) that φ: at no situation s', compatible with ('accessible from') the situation s, it is the case that φ (either s' is incomplete with respect to φ or $\neg\varphi$ holds there).

$$s \Vdash \neg\varphi \quad \text{iff} \quad (\forall s' \in S)(sCs' \text{ implies } s' \nVdash \varphi)$$

Informally speaking, the agent can explicitly deny some hypothesis (a piece of data) only if no research group in her neighborhood claims it is true. This condition also has a normative side: she has to be skeptical in the sense that she denies everything not positively supported by any of her colleagues (in the situations related to her actual situation).

If we want to grant negative facts the same basic level as positive facts, we can read the clause for the definition of compatibility in the other direction: the agent can relate her actual situation just to the situations which do not contradict her negative facts.

Depending on the properties of the compatibility relation we obtain different kinds of negations. We shall shortly comment on them.

The compatibility relation is in general not reflexive: inconsistent situations are not self-compatible and so reflexivity holds only for consistent situations. It is clear that for an inconsistent self-compatible situation the clause for negation would not work. On the other hand, inconsistent situations can be compatible with some incomplete situations.

Nor is C transitive. Let us have situations x, y, z such that $x \Vdash \varphi$, $z \Vdash \neg\varphi$, and y does not include either φ or $\neg\varphi$. Assume that xCy and yCz. Then according to the definition of negation it cannot be that xCz.

It is quite reasonable to assume that C is *symmetric*. This condition implies that we get only one negation (otherwise we would get left and right negation) and we get the 'unproblematic' half of the law of double negation (if $x \Vdash \varphi$, then $x \Vdash \neg\neg\varphi$).

We also assume C is *directed* and *convergent*. Directedness means that there is at least one compatible situation for each $x \in S$. Convergence says that there is a maximal compatible situation $x^\star$. (See appendix A.)

Maximal compatible situations (with respect to x) can be inconsistent about everything not considered in x. From the symmetry of C we obtain $x \unlhd x^{\star\star}$. If we assume, moreover, $x \unrhd x^{\star\star}$, then we get the operation $\star$ with the property $x = x^{\star\star}$, i.e. the *Routley star*. The definition of negation-validity is then written in the form:

$$x \Vdash \neg\varphi \quad \text{iff} \quad x^\star \nVdash \varphi$$

The Routley star has been one of the controversial points of the Routley–Meyer semantics, but in our motivation it has a quite natural explanation: if compatible situations represent colleagues from different research groups our agent collaborates with, then the maximal compatible situation correspond to a colleague ('boss') who has all the information the other colleagues from the group have. Then if the agent wants to accept some negative clause she does not have to speak to each of the colleagues and ask his/her opinion, she just asks the 'boss' directly and knows that bosses opinion represents the opinions of the entire compatible research group.

This completes our exposition of relational semantics for relevant logics. We now move to epistemic modalities.

3 Knowledge in relevant framework

There have been some attempts to combine an epistemic and relevant framework (see (Cheng, 2000) and (Wansing, 2002)), but they have a different aim then our approach.

From a purely technical point of view there are a number of ways to introduce modalities in the relevant framework — Greg Restall in (Restall, 2000) provides a nice general overview. As we mentioned, the relevant framework already contains modal notions. We therefore decided to use these notions to introduce epistemic modalities rather than to introduce new ones.

In the classical epistemic frame what an agent knows in a world w is defined as what is true in all epistemic alternatives of w, which are given by the corresponding accessibility relation. Our idea of the agent as a scientist processing some kind of data requires a different approach.

We assume our agent in her current situation s observes (has a direct approach to) some data, represented by formulas which are true at s. She is aware of the fact that these data might be unreliable (or even inconsistent). In order to accept some of the current data as knowledge the agent requires a confirmation from some 'independent' resources.

In our approach resources are situations dealing with the same kind of data available in the current situation. A resource shall be more elementary than the current situation, i.e., it should not contain more data (a resource is below s in the $\trianglelefteq$-relation). Also the data from the resource should not contradict the data in the current situation (a resource is compatible with s).

Definition 1 (Knowledge)**.**

$$s \Vdash K\varphi \quad \text{iff} \quad (\exists x)(sC^{\lhd}x \text{ and } x \Vdash \varphi),$$

where $sC^{\lhd}x$ iff sCx and $x \trianglelefteq s$ and $x \neq s$.

In short, φ is known iff there is a resource ('lower' compatible situation different from the actual one) validating φ.

We allowed our agent to deal with inconsistent data in order to get a more realistic picture. However, the agent should be able to separate inconsistent data. The modality we introduced provides us with such an appropriate filter. Let us assume both φ and $\neg\varphi$ are in s (e.g., our agent might receive such inconsistent information from two different sources). The agent considers both φ and $\neg\varphi$ to be possible, but neither of them is confirmed information as according to the definition, no situation compatible to s can contain either φ or $\neg\varphi$.

Basic properties

It is to be expected that our system blocks all the undesirable properties of both material and strict implication. Moreover, we ruled out the validity of some of the properties of 'classical' epistemic logics that we have criticized, in particular, both positive and negative introspection, as well as some closure properties.

Let us have a relevant frame $\mathbf{F} = \langle S, L, C, \trianglelefteq, R\rangle$. Recall that the truth in the frame $\mathbf{F}$ corresponds to the truth in the logical situations of $\mathbf{F}$ (under any valuation). We will also use the stronger notion of truth in all situations of $\mathbf{F}$ (under any valuation). From the viewpoint of our motivation the latter notion is more interesting as our agent might happen to be in other situations than the logical ones.

Our approach makes the 'truth axiom' **T** valid. For any situation $s \in S$, if φ is known at s ($s \Vdash K\varphi$), then there is a $\trianglelefteq$-lower compatible witness with φ true, which makes φ to be true at s as well. Thus, formula

$$K\alpha \rightarrow \alpha$$

is valid.

The axiom **K** and the necessity rule, common to all normal epistemic logics, fail. First, let us assume that φ is valid formula. The necessity rule

($\dfrac{\varphi}{K\varphi}$) would imply the validity of $K\varphi$. $\models \varphi$ means that φ is true in every logical situation l. However, for $l \Vdash K\varphi$ a confirmation from a different resource is required, there must be a situation x such that $x \Vdash \varphi$ and $lC^{\lhd}x$, which, in general, does not need to be the case.

Second, in our interpretation the validity of axiom **K** is not well motivated and does not hold. **K** is in fact a 'distribution of confirmation': If an implication is confirmed then the confirmation of the antecedent implies the confirmation of the consequent.

$$\not\models K(\alpha \to \beta) \to (K\alpha \to K\beta)$$

Introspection As we defined knowledge as independently confirmed data, the epistemic axioms **4** and **5** correspond in our framework to a 'second order confirmation' rather than to introspection. It is easy to see that both axioms fail.

$$\not\models K\alpha \to KK\alpha,$$
$$\not\models \neg K\alpha \to K\neg K\alpha$$

Necessity and possibility We do not introduce possibility using the standard definition $M\varphi \overset{\text{def}}{\equiv} \neg K\neg\varphi$. Our idea of epistemic possibility is that our agent considers all the data available at the current situation as possible. If we introduce formally $s \Vdash M\varphi$ as $s \Vdash \varphi$, then it follows from the **T** axiom that in all situations necessity implies possibility:

$$(\forall s \in S)(s \Vdash K\varphi \to M\varphi)$$

However for the standard dual possibility this is not true.

$$\not\models K\varphi \to \neg K\neg\varphi$$

Let us comment on the relation of negation and necessity in our framework. There is a difference between $s \nVdash K\varphi$ and $s \Vdash \neg K\varphi$. The former simply says that φ is not confirmed at the current situation s, while the latter says that φ is not confirmed in the situations compatible with s. From this point of view it is uncontroversial that both $K\varphi$ (confirmation in the current situation) and $\neg K\varphi$ (the lack of confirmation in the compatible situations) might be true in some situation s (the necessary condition is that s is not compatible with itself).

Closure properties It is easy to see that the modal Modus Ponens

$$\frac{K\alpha \quad K(\alpha \to \beta)}{K\beta}$$

does not hold (for the reasons given in the section on **K** axiom). However, its weaker version

$$\frac{K\alpha \quad K(\alpha \to \beta)}{\beta}$$

holds not only in logical situations, but in all situations. If $K\alpha$ and $K(\alpha \to \beta)$ are true in any $s \in S$, then $s \Vdash \beta$. Axiom **T** and the assumption $Rsss$ are crucial here.

Contradiction in our system is non-explosive: φ and $\neg\varphi$ might hold in a contradictory situation, which need not be connected to any situation where ψ holds.

$$\not\models (\varphi \wedge \neg\varphi) \to \psi$$

On the other hand, the knowledge of contradiction implies anything (as a contradiction is never confirmed):

$$\models K(\varphi \wedge \neg\varphi) \to \psi$$

Modal adjunction also does not hold — if $K\alpha$ and $K\beta$ are true in s, then obviously $(\alpha \wedge \beta)$ is true there because of the truth axiom but $K(\alpha \wedge \beta)$ does not need to be true in s. (If each of α and β is confirmed by some resource, there still might be no resource confirming their conjunction.)

4 Conclusion

We introduced a system of epistemic logic based on the framework of relevant logic. We gave an epistemic interpretation of the relational semantics for relevant logics and defined epistemic modalities motivated by this interpretation. Instead of introducing additional relations into the framework, we argued in favor of using modalities based on the relations already contained in the frame.

The whole project is at an initial stage: there is much to be done both technically and in the area of interpretation. In particular we shall develop in a more detail the epistemic interpretation of our framework, give an axiomatization of our system, and characterize its formal properties.

A Relevant logic R

There are more formal systems that can be called relevant logic. From the proof-theoretical viewpoint, all of them are considered to be substructural logics (see (Restall, 2000) and (Paoli, 2002)). Here we present the axiom system and (Routley–Meyer) semantics from (Mares, 2004) with some elements from (Restall, 1999).

Syntax

We use the language of classical propositional logic with signs for atomic formulas $\mathcal{P} = \{p, q, \dots\}$, formulas being defined in the usual way:

$$\varphi ::= p \mid \neg\psi \mid \psi_1 \vee \psi_2 \mid \psi_1 \wedge \psi_2 \mid \psi_1 \rightarrow \psi_2$$

Axiom schemes

1. $A \rightarrow A$
2. $(A \rightarrow B) \rightarrow ((B \rightarrow C) \rightarrow (A \rightarrow C))$
3. $A \rightarrow ((A \rightarrow B) \rightarrow B)$
4. $(A \rightarrow (A \rightarrow B)) \rightarrow (A \rightarrow B)$
5. $(A \wedge B) \rightarrow A$
6. $(A \wedge B) \rightarrow B$
7. $A \rightarrow (A \vee B)$
8. $B \rightarrow (A \vee B)$
9. $((A \rightarrow B) \wedge (A \rightarrow C)) \rightarrow (A \rightarrow (B \wedge C))$
10. $(A \wedge (B \vee C)) \rightarrow ((A \wedge B) \vee (A \wedge C))$
11. $\neg\neg A \rightarrow A$
12. $(A \rightarrow \neg B) \rightarrow (B \rightarrow \neg A)$

Strong logical constants $\otimes$ (group conjunction, fusion) and $\oplus$ (group disjunction) are definable by implication and negation:

- $(A \oplus B) \stackrel{\text{def}}{\equiv} \neg(\neg A \rightarrow B)$
- $(A \otimes B) \stackrel{\text{def}}{\equiv} \neg(\neg A \oplus \neg B)$

Rules

Adjunction From A and B infer $A \wedge B$.

Modus Ponens From A and $A \rightarrow B$ infer B.

Routley–Meyer semantics

An R-frame is a quintuple $\mathbf{F} = \langle S, L, C, \trianglelefteq, R\rangle$, where S is a non-empty set of situations and $L \subseteq S$ is a non-empty set of logical situations. The relations $C \subseteq S^2$, $\trianglelefteq \subseteq S^2$, and $R \subseteq S^3$ were introduced in section 2, here we sum up their properties.

Properties of the relation R The basic property of R:

$$\text{if } Rxyz,\ x' \trianglelefteq x,\ y' \trianglelefteq y, \text{ and } z \trianglelefteq z', \quad \text{then } Rx'y'z'.$$

This means that the relation R is monotonic with respect to the involvement relation.

Moreover it is required that:

(r1) $Rxyz$ implies $Ryxz$;

(r2) $R^2(xy)zw$ implies $R^2(xz)yw$, where R^2xyzw iff $(\exists s)(Rxys$ and $Rszw)$;

(r3) $Rxxx$;

(r4) $Rxyz$ implies $Rxz^\star y^\star$.

Properties of the relation C Compatibility between two states is inherited by the states involved in them ('less informative states'):

$$\text{If } xCy,\ x_1 \trianglelefteq x, \text{ and } y_1 \trianglelefteq y, \text{ then } x_1Cy_1.$$

Moreover, we require the following properties:

(c1) (symmetricity) xCy implies yCx;

(c2) (directedness) $(\forall x)(\exists y)(xCy)$;

(c3) (convergence) $(\forall x)(\exists y(xCy)$ implies $(\exists x^\star)(xCx^\star$ and $\forall z(xCz$ implies $z \trianglelefteq x^\star)))$;

(c4) $x \trianglelefteq y$ implies $y^\star \trianglelefteq x^\star$;

(c5) $x^{\star\star} \trianglelefteq x$.

Model R-model $\mathbf{M}$ is a R-frame $\mathbf{F}$ with a valuation function $v\colon \mathcal{P} \to 2^S$. The truth of a formula at a situation is defined in the following way:

- $s \Vdash p$ iff $s \in v(p)$,
- $s \Vdash \neg\varphi$ iff $s^\star \nVdash \varphi$,

- $s \Vdash \psi \wedge \varphi$ iff $s \Vdash \psi$ and $s \Vdash \varphi$,
- $s \Vdash \psi \vee \varphi$ iff $s \Vdash \psi$ or $s \Vdash \varphi$,
- $s \Vdash (\varphi \rightarrow \psi)$ iff $(\forall y, z)(Rsyz$ implies $(y \Vdash \varphi$ implies $z \Vdash \psi))$.

As we already said, the truth of a formula in a model and in a frame, respectively, is defined as truth in all logical situations of this model/frame. As usual, R-tautologies are formulas true in all relevant frames. Whenever φ is a R-tautology, we write $\models \varphi$ and say that φ is a valid formula.

The condition (r1) validates the implicative version of *Modus Ponens* (axiom schema 3). It does not validate the conjunctive version $(A \wedge (A \rightarrow B)) \rightarrow B$, which requires (r3).

(r2) corresponds to the '*exchange rule*' $(A \rightarrow (B \rightarrow C)) \rightarrow (B \rightarrow (A \rightarrow C))$, which is derivable from the axioms given above.

(r4) validates *contraposition* (axiom schema 12). If we work without the Routley star, this can be rewritten as:

$$Rxyz \text{ implies } (\forall z' Cz)(\exists y' Cy)(Rxy'z').$$

Directedness and convergence conditions are necessary for the definition of the Routley star. From (c1) we obtain the validity of $(A \rightarrow \neg\neg A)$ and from the last condition (c5) we get the axiom schema 11.

Ondrej Majer
Institute of Philosophy, Academy of Sciences of the Czech Republic
Jilská 1, 110 00 Praha 1
`majer@site.cas.cz`
`http://logika.flu.cas.cz`

Michal Peliš
Institute of Philosophy, Academy of Sciences of the Czech Republic
Jilská 1, 110 00 Praha 1
`pelis@ff.cuni.cz`
`http://logika.flu.cas.cz`

References

Cheng, J. (2000). A strong relevant logic model of epistemic processes in scientific discovery. In E. Kawaguchi, H. Kangassalo, H. Jaakkola, & I. Hamid (Eds.), *Information modelling and knowledge bases XI* (pp. 136–159). Amsterdam: IOS Press.

Duc, H. N. (2001). *Resource-bounded reasoning about knowledge.* Unpublished doctoral dissertation, Faculty of Mathematics and Informatics, University of Leipzig.

Fagin, R., Halpern, J., Moses, Y., & Vardi, M. (2003). *Reasoning about knowledge.* Cambridge, MA: MIT Press.

Mares, E. (2004). *Relevant logic.* Cambridge: Cambridge University Press.

Mares, E., & Meyer, R. (1993). The semantics of **R4**. *Journal of Philosophical Logic*, *22*, 95–110.

Paoli, F. (2002). *Substructural logics: A primer.* Dordrecht: Kluwer.

Restall, G. (1993). Simplifeid semantics for relevant logics (and some of their rivals). *Journal of Philosophical Logic*, *22*, 481–511.

Restall, G. (1995). Four-valued semantics for relevant logics (and some of their rivals). *Journal of Philosophical Logic*, *24*, 139–160.

Restall, G. (1996). Information flow and relevant logics. In *Logic, language and computation: The 1994 Moraga proceedings. CSLI Lecture Notes* (Vol. 58, pp. 463–477). Stanford, CA: CSLI.

Restall, G. (1999). Negation in relevant logics: How I stopped worrying and learned to love the Routley star. In D. Gabbay & H. Wansing (Eds.), *What is negation?* (Vol. 13, pp. 53–76). Dordrecht: Kluwer.

Restall, G. (2000). *An introduction to substructural logics.* London–New York: Routledge.

Wansing, H. (2002). Diamonds are a philosopher's best friends. *Journal of Philosophical Logic*, *31*, 591–612.