

On the Relation between the Coherence and Foundations Theories of Belief Revision

Alvaro del Val
Robotics Lab
Computer Science Department
Stanford University
Stanford, CA 94305
delval@cs.stanford.edu

Abstract

Two recent papers, (Gärdenfors 1990; Doyle 1992), try to assess the relative merits of the two main approaches to belief revision, the foundations and coherence theories, but leave open the question of the mathematical connections between them. We answer this question by showing that the foundations and coherence theories of belief revision are mathematically equivalent. The result also has consequences for non-monotonic reasoning, as it entails that Poole's system of default reasoning and Shoham's preferential logic are expressively equivalent, in that they can represent the same set of non monotonic consequence relations.

Introduction

Two major approaches to belief revision can be distinguished, according to the role assigned in the belief revision process to the agent's reasons for holding his or her beliefs. In the *foundations theory* of belief revision, the agent's beliefs are seen as having a structure beyond the purely logical relations among them. In particular, certain beliefs are *justified* by some other beliefs, which in turn might be justified by still other beliefs, etc., with a distinguished set of "basic" or "self-justified" beliefs providing the foundation for the whole edifice. When the agent's beliefs are to be revised, some of these basic beliefs might have to be retracted; as a result some other beliefs will become unjustified, and according to the foundations approach they should be retracted as well. In contrast, in the *coherence theory* of revision the goal is only to maintain the overall consistency of the agent's beliefs, while retracting as few beliefs as possible during revision. In the basic approach, all beliefs are in principle accorded the same status, and the agent will keep a belief whenever he or she can consistently do so, even when the original reasons for holding that belief are retracted.

Two recent papers, (Gärdenfors 1990; Doyle 1992), try to assess the relative merits of each approach. Recognizing that the question is unlikely to be solved by informal arguments, both authors consider the question of the mathematical connections between the approaches. Specifically, Gärdenfors tries to show that in

many cases the notion of a "reason for belief," which seems to be fundamental to the foundational approach, can be "reconstructed" from a coherentist point of view, specifically using the notion of "epistemic entrenchment," Gärdenfors' preferred way to conceptualize the coherence approach. The proposal is only suggestive, as Gärdenfors admits it is formally flawed. Doyle expands on the flaws of this proposal, and suggests that it should be possible to encode coherence revision operators in a foundational framework.

Thus, the question of the mathematical connections between both approaches remains open. We answer this question in this paper, by showing that the foundations theory is equivalent to the coherence theory. Answering this question of course requires to be more precise about the formal definition of both approaches. Whereas we will, with Doyle and Gärdenfors, take AGM-like revision as our model of coherentist revision, we depart from both of them in our choice of a formal model for the foundations theory. In particular, we follow Nebel (1991) in that the notion of "reasons for beliefs" will play no explicit formal role in our definition of the foundational approach, where we will depart only slightly from the "syntax-based approach" advocated by this and other authors.

In the next two sections, we present the formal model we use for the coherence and foundations theory, respectively. We then present the main technical results of the paper, and finish by discussing the implications of these results on the expressiveness of two non-monotonic frameworks.

In the rest of the paper, we assume a propositional language \mathcal{L} obtained by closing a finite set of symbols \mathcal{P} under the usual boolean connectives. \mathcal{W} is the set of all interpretations of \mathcal{L} . $Mod(\psi)$, for any $\psi \in \mathcal{L}$, denotes the set of models of ψ . \vdash stands for propositional consequence, and for any $\Sigma \subseteq \mathcal{L}$, $Cn(\Sigma)$ is $\{\varphi \mid \Sigma \vdash \varphi\}$ the logical closure of Σ . A preorder is a reflexive and transitive relation. For any preorder \leq and any subdomain S of \leq , $Min(S, \leq)$ denotes the set of minimal elements of S under \leq .

Finally, the following notation will be useful. Let

$$\Sigma \Downarrow \neg \mu = \{\Gamma \subseteq \Sigma \mid \Gamma \not\vdash \mu \text{ and } \forall \Theta \subseteq \Sigma, \text{ if } \Gamma \subset \Theta \text{ then } \Theta \vdash \mu\}$$

be the set of maximal subsets of Σ that do *not* entail μ . The set $\Sigma \Downarrow \mu$ can also be filtered by incorporating a “preference preorder” \preceq over subsets of Σ , whose strict part is written \prec , defining $\Sigma \Downarrow \mu = \{\Gamma \subseteq \Sigma \mid \Gamma \not\vdash \mu \text{ and } \forall \Theta \subseteq \Sigma, \text{ if } \Theta \prec \Gamma \text{ then } \Theta \vdash \mu\}$. We require \preceq to extend set containment, *i.e.* to satisfy $\Theta \prec \Gamma$ whenever $\Gamma \subset \Theta$.

The coherence theory

The main formal representative of the coherence theory of belief revision is the theory developed in (Alchourrón, Gärdenfors, & Makinson 1985; Gärdenfors 1988). The AGM approach to revision has become identified to a great extent with its *normative* side: the authors put forward a set of “postulates” that, they claim, any “rational” revision operator should satisfy. For readability, and since we are considering only the finitary propositional case, we follow the presentation of (Katsuno & Mendelzon 1991). Using \circ to denote a revision operator, the postulates are:

- (R1) $\psi \circ \mu$ implies μ .
- (R2) If $\psi \wedge \mu$ is satisfiable then $\psi \circ \mu$ is equivalent to $\psi \wedge \mu$.
- (R3) If μ is satisfiable then $\psi \circ \mu$ is also satisfiable.
- (R4) If $\models \psi_1 \equiv \psi_2$ and $\models \mu_1 \equiv \mu_2$ then $\psi_1 \circ \mu_1$ is equivalent to $\psi_2 \circ \mu_2$.
- (R5) $(\psi \circ \mu) \wedge \phi$ implies $\psi \circ (\mu \wedge \phi)$.
- (R6) If $(\psi \circ \mu) \wedge \phi$ is satisfiable then $\psi \circ (\mu \wedge \phi)$ implies $(\psi \circ \mu) \wedge \phi$.

We will also consider the following two postulates, from (Katsuno & Mendelzon 1991), as a weaker alternative to (R6):

- (R7) If $\psi \circ \mu_1$ implies μ_2 and $\psi \circ \mu_2$ implies μ_1 then $\psi \circ \mu_1$ is equivalent to $\psi \circ \mu_2$.
- (R8) $(\psi \circ \mu_1) \wedge (\psi \circ \mu_2)$ implies $\psi \circ (\mu_1 \vee \mu_2)$.

Definition 1 A coherence revision operator is any operator \circ satisfying postulates (R1)–(R5), (R7), and (R8). An AGM operator is a coherence operator which in addition satisfies (R6).

Coherence revision operators can be characterized by the following representation theorem, due to (Katsuno & Mendelzon 1991). A *revision assignment* is a function that assigns to each formula ψ a binary relation \leq_ψ over \mathcal{W} , the set of all interpretations of \mathcal{L} . The revision assignment is said to be *faithful* iff:

1. $\text{Min}(\mathcal{W}, \leq_\psi) = \text{Mod}(\psi)$ for any satisfiable ψ ; and
2. $\leq_\psi = \leq_\phi$ whenever $\models \psi \equiv \phi$.

Theorem 1 A revision operator \circ satisfies conditions (R1)–(R5), (R7) and (R8) (respectively, (R1)–(R6)) iff there exists a faithful revision assignment that maps each formula ψ to a partial (respectively total) preorder \leq_ψ such that:

$$\text{Mod}(\psi \circ \mu) = \text{Min}(\text{Mod}(\mu), \leq_\psi)$$

Why should these postulates be regarded as characterizing a *coherence* theory of belief revision? The main reason has to do, in our view, with the first representation theorem used to characterize operators satisfying them, which is different from the one just given. According to the coherence theory, as said, the main criterion in deciding whether to preserve certain beliefs in the face of revision is whether they can be consistently held after the new information is incorporated, in which case they should be preserved. Thus, a natural way to capture this idea is to view revision as a two step process. In the first step, the agent checks whether the new information is consistent with his or her beliefs, and, if this is not the case, withdraws as few beliefs as possible so as to restore consistency; in the second step, the beliefs kept in the previous stage are conjoined with the new information. (The two steps correspond, respectively, to the AGM operations of contraction and expansion.) As we will see, this is very similar to the approach taken by the foundations theory of revision, with the only difference that the latter considers only a distinguished set of basic beliefs. In the coherence theory, in contrast, all beliefs are, at least in principle, accorded the same status.

Formally, the idea of removing as few beliefs as possible in the first step can be captured in terms of the \Downarrow notation introduced in the first section. Suppose the agent’s beliefs are (finitely) represented by some sentence ψ , to be revised with some new information μ . Under most, though by not means all, formal conceptions of belief, the agent will also believe in any logical consequence of ψ , and thus the agent’s beliefs are given by $Cn(\psi)$. Because the coherence theory accords all beliefs, in principle, the same status, all beliefs must be considered in minimizing retracted beliefs. This means that the first step in revision should be captured in terms of the set $Cn(\psi) \Downarrow \neg\mu$. If this set is a singleton, say $\{\Gamma\}$, the second step can be captured by defining the result of revising ψ with μ to be $Cn(\Gamma \cup \{\mu\})$. Otherwise, there is some choice as to what to do, e.g. choosing one, some, or all the elements of $Cn(\psi) \Downarrow \neg\mu$. Abstracting away from the details, we can simply assume that there is a selection function $S_\psi : \mathcal{P}(Cn(\psi)) \rightarrow \mathcal{P}(Cn(\psi))$, satisfying $\emptyset \subset S_\psi(\Psi) \subseteq \Psi$, and define revision by:

$$Cn(\psi \circ \mu) = \bigcap_{\Gamma \in S_\psi(Cn(\psi) \Downarrow \neg\mu)} Cn(\Gamma \cup \{\mu\}).$$

The selection function can be seen as expressing some preferences on the agent’s beliefs. It turns out that, by placing certain conditions on this function, the class of AGM operators can be fully characterized by means of a representation theorem (Alchourrón, Gärdenfors, & Makinson 1985), a theorem that legitimates the identification of the AGM approach with the coherence theory.

The foundations theory

Foundational approaches, as said, postulate a distinction between “basic” or self-justifying beliefs and other beliefs, which should be ultimately justified in terms of the former. As an example, suppose we initially believe that some particular animal is a mammal and that every mammal has lungs; then we will also believe that the animal has lungs. We can represent these beliefs with the database $\{m, m \supset l, l\}$. If we are now told that the animal is not a mammal after all, we need to revise our beliefs with $\neg m$. Many coherence revision operators, such as *e.g.* the one proposed by (Dalal 1988), would yield a revised database equivalent to $\neg m \wedge l$; *i.e.* we would retain the belief that the animal has lungs, even if we no longer have any reason to believe it. But, one could argue, this would only be warranted if this belief did not “depend” on the belief that it is a mammal (say, we have independently observed that it has lungs). Coherence approaches appear *prima facie* ill-suited to make this kind of distinction; in a foundational approach, in contrast, the first case would correspond to treating $\{m, m \supset l\}$ as the set of self-justifying beliefs, and the second case to treating $\{m, l\}$ as basic beliefs. Assuming that we want to preserve as many basic beliefs as possible, a typical foundational approach would revise each database differently, yielding respectively $\{m \supset l, \neg m\}$ and $\{\neg m, l\}$ as revised databases. We thus capture the distinction between having an independent *reason* for believing that the animal has lungs and not having it, even though the notion of “reasons for belief” plays no explicit formal role.

One straightforward way to capture this distinction between basic and non-basic beliefs is to base it on the proof-theoretic notion of derivability. Given a finite axiomatization of a theory, we can take the axioms to be the basic beliefs; any other beliefs about the domain should be justified, *i.e.* *derivable* from the axioms in the underlying logic. This is the intuition behind syntax-based approaches to revision, the main representatives, in our view, of the foundational theory — see *e.g.* (Fagin, Ullman, & Vardi 1983; Makinson 1985; Ginsberg 1986; Nebel 1989; Benferhat *et al.* 1993).

Thus, if the finite set of sentences Ψ axiomatizes the agent’s beliefs, and we take this set to be identical with the set of basic beliefs, syntax-based revision can be defined by $Cn(\Psi \circ \mu) = \bigcap_{\Gamma \in \Psi \downarrow \neg \mu} Cn(\Gamma \cup \{\mu\})$, where $\downarrow = \Downarrow$ when $\preceq = \supseteq$, and is otherwise based on some preorder \preceq over subsets of Ψ , or more generally over finite subsets of \mathcal{L} , that extends set containment.

Though the syntax-based approach appears to capture some of the intuitions behind the foundational approach, it also has a serious drawback, namely, the result of revision becomes extremely dependent on the syntactic form of the database. Even apparently meaningless distinctions such as that between $\{a, b\}$ and $\{a \wedge b\}$ have an effect in revision. This has two highly undesirable consequences. First, it forces the user

to write down the database so as to reflect the distinction between basic and derived beliefs, relegating any other concerns such as conciseness, understandability, or suitability for efficient inference, which are equally important in choosing the “right” axiomatization. Second, even assuming that the user is willing to write down the axioms with such an exclusive concern for proper revision behavior, any further modification of the database, such as caching the results of inferences and storing them back with the database, or applying equivalence-preserving transformations to the database for optimization purposes,¹ will also affect the results of revision.

This is clearly unsatisfactory. What is needed is the ability to separate the specification of the basic beliefs from the axiomatization of the database. For this reason, we will depart slightly from traditional presentations of syntax-based approaches, by introducing a *basic beliefs function* $\Sigma : \mathcal{L} \rightarrow \mathcal{P}(\mathcal{L})$ that maps any formula ψ into an associated set of basic beliefs $\Sigma(\psi)$ (also written Σ_ψ for brevity). We require Σ to satisfy:

1. $\Sigma(\psi) = \Sigma(\varphi)$ whenever $\vdash \psi \equiv \varphi$.
2. $Cn(\Sigma(\psi)) = Cn(\psi)$.
3. $\Sigma(\psi)$ is finite.

Definition 2 A foundations revision operator is an operator \circ defined by

$$Cn(\psi \circ \mu) = \bigcap_{\Gamma \in \Sigma_\psi \downarrow \neg \mu} Cn(\Gamma \cup \{\mu\}),$$

where Σ is a basic beliefs function, and \downarrow is as defined above, in terms of some preorder \preceq that extends set containment. Furthermore, a foundations operator is:

- a total preorder operator iff \preceq is a total preorder;
- a basic foundations operator iff $\Downarrow = \downarrow$, *i.e.* $\preceq = \supseteq$;
- a TO foundations operator iff it is basic and $\Sigma_\psi = \{\sigma_1, \dots, \sigma_n\}$, for some n , where $\sigma_{i+1} \in Cn(\sigma_i)$ for $1 \leq i < n$.

As we will see later, the class of foundations operators and the class of basic foundations operators are identical. TO operators are included because they can capture the class of AGM coherence operators. If we ignore questions of non-deterministic revision (on which more below), the class of TO operators is identical to the total preorder class. It is also equivalent to the total preorder operator that Nebel (1991) calls “unambiguous prioritized revision.”

The use of a basic beliefs function is analogous to the use of a preference preorder on models in coherence operators as per theorem 1, *i.e.* is a device for specifying a belief revision policy. Note that we do introduce a principle of syntax independence in the definition of this function, and thus of foundations operators. In

¹This includes *e.g.* transformation to CNF. Note also that syntax-dependence does not go away by restricting the syntactic form of the database, *e.g.* clausal, Horn, etc.

our view, arguments for syntax-dependence boil down to the practical convenience of using the database axiomatization as a device for specifying a foundational revision policy, that is, the set of basic beliefs. This convenience is unaffected by our reformulation, since in practice one is given a single initial database; we can take its axioms as basic beliefs, and stipulate that any equivalent database has the same associated basic beliefs as the ones given. Computationally, therefore, we can proceed exactly as in the syntax-based approach, since the proof theoretic approach characteristic of syntax-based revision is preserved (with basic beliefs replacing the axioms).²

As for the second condition on Σ , the requirement that all basic beliefs are believed (that $Cn(\Sigma_\psi) \subseteq Cn(\psi)$) is obvious. To see that the converse inclusion is also needed, note that $Cn(\psi \circ true) = Cn(\Sigma_\psi)$, and thus if $Cn(\Sigma_\psi) \subset Cn(\psi)$ some beliefs would be lost when revising with a tautology. The third requirement, in conjunction with the second, entails that Σ_ψ provides a finite axiomatization of ψ .

This concludes the presentation of the formal model of foundational belief revision that we use. As said, both Gärdenfors and Doyle take the JTMS (justification based truth maintenance system, (Doyle 1979)) as the paradigm of foundational revision. Why do we choose a different model? We have argued, following (Nebel 1991), that the distinction between basic and non-basic beliefs is the only essential aspect of the foundations theory, without any need for an explicit concept of justification. In our view, the role of a TMS (not just a JTMS) in the context of belief revision is simply to allow us to easily detect whether the new formula contradicts previous beliefs and to trace back the basic beliefs underlying this contradiction, by caching inferences as well as the reasons for beliefs. Thus, from the point of view of revision the role of a TMS is simply to facilitate the computation of $\Sigma \downarrow \neg\mu$ (see (Benferhat *et al.* 1993) for a detailed treatment of this topic).

The specific choice of a JTMS as the paradigm of foundational revision presents two additional problems. First, inference with the JTMS is equivalent to inference in general logic programs, with the semantics of autoepistemic or default logic (Pimentel & Rodi 1991; Elkan 1990; Reinfrank, Dressler, & Brewka 1989); while the problem of revision is well defined for any logic, we see no reason to take revision in a non-classical logic as a paradigm of foundational revision. Second, all the revision procedures proposed for the JTMS (Doyle 1979; Elkan 1990; Pimentel & Rodi 1991) share a fundamental limitation, which in our view makes them inadequate as *general* models of revision. Namely, when a contradiction is de-

²A potential problem arises for iterated revisions, if the agent ends up having the same beliefs at a later time, but for different reasons (different basic beliefs). This is easily solved by time indexing Σ , allowing the basic beliefs to vary with time. See (del Val & Shoham 1994).

tected, revision has to be performed by retracting literals (nodes) that appear as non monotonic antecedents in JTMS's justifications; but when this is not possible the system will remain in an inconsistent state.

The equivalence of both approaches

In this section we show that the coherence and foundational approaches are equivalent. The first direction, from foundational to coherence operators, is easy, and can be found in a less general form in the literature.

Theorem 2 *For any foundations revision operator \circ_F there exists a coherence revision operator \circ_C such that $Mod(\psi \circ_F \mu) = Mod(\psi \circ_C \mu)$.*

Theorem 3 *For any total preorder or TO foundations revision operator \circ_F there exists an AGM coherence revision operator \circ_C such that $Mod(\psi \circ_F \mu) = Mod(\psi \circ_C \mu)$.*

These two theorems capture a much wider family of syntax-based operators satisfying the respective set of postulates than those considered by Nebel (1989; 1991). For example, he introduces “prioritized revision,” in which the formulas of the database are partitioned into a set of totally ordered “priority strata,” defining a lexicographic “prioritized” ordering on subsets of the database, and shows that it satisfies postulates (R1)–(R5). It is easy to see however from theorem 2 that if the ordering on the strata is allowed to be partial (in the style of (Grosz 1991)) these postulates are still satisfied, together with (R7) and (R8). Similarly, Nebel introduces “unambiguous prioritized revision,” in which each strata is a singleton, showing that it satisfies the AGM postulates. It is a consequence of theorem 3, for example, that the operators introduced in (Ginsberg 1986) based on “modular orders” also satisfy these postulates, since modular orders can be easily mapped into total preorders. Our characterizations can also accommodate a variety of “voting schemes,” as suggested in (Doyle 1991), as well as all the operators proposed in (Benferhat *et al.* 1993).

The most novel contribution of this paper is however given by the converses of the previous two theorems:

Theorem 4 *For any coherence revision operator \circ_C there exists a basic foundations revision operator \circ_F such that for every ψ and μ , $Mod(\psi \circ_C \mu) = Mod(\psi \circ_F \mu)$.*

Theorem 5 *For any AGM coherence revision operator \circ_C there exists a basic foundations revision operator \circ_F satisfying the TO condition and such that for every ψ and μ , $Mod(\psi \circ_C \mu) = Mod(\psi \circ_F \mu)$.*

It follows that, at least for finitary propositional languages, there is no choice to be made between the coherence and foundations theory of belief revision: they are mathematically equivalent. Theorems 2 and 3 show that any foundations revision operator can be seen as a coherence operator, in essence exploiting the fact that

coherence operators allow for using preferences over beliefs in determining the revised database. Theorems 4 and 5, in turn, show that it is always possible to choose the set of basic beliefs so as to encode any coherence revision operator as a foundational operator.³

Note also that by theorem 2, any foundations operator satisfies (R1)–(R5) and (R7) and (R8), and that by theorem 4, any operator satisfying these postulates can be defined by means of a *basic* foundations operator. It follows that the two classes of foundations operators are identical. Similarly, the classes of TO and total preorder foundations operators are identical.

We omit proofs for lack of space. The easiest way to explain the connection between AGM and TO operators is probably to note that the set of sets of models $\{Mod(\sigma_1), \dots, Mod(\sigma_n)\}$ forms an “embedded system of spheres,” in the sense of Grove’s (1988) representation theorem for AGM revision. Establishing the connection for non-AGM coherence operators (that is, operators based on a partial preorder on models) requires more work, based on a similar construction. The payoff for this additional work will be apparent in the next section. The results can be extended to the infinitary case as long as we can assume that the equivalence classes derivable from the preorders defining coherence operators are finitely axiomatizable.

We end this section by noting the following fact:

Theorem 6 *Let \circ be a foundational operator, and suppose $\Sigma_\psi \downarrow \neg\mu$ is a singleton for every ψ and μ . Then \circ satisfies (R6).*

The practical utility of having $\Sigma_\psi \downarrow \neg\mu$ be a singleton was already noted in (Nebel 1989), and it becomes clearer when we consider the non-deterministic syntax-based approach advocated in (Fagin *et al.* 1986), see also (Doyle 1991). In this approach, each element of $\Psi \downarrow \neg\mu$, where Ψ is a finite sets of formulas, is taken to generate an *alternative* revised database; $\Psi \circ \mu$ is taken to be a set of databases, namely the set $\{\Gamma \cup \{\mu\} \mid \Gamma \in \Psi \downarrow \neg\mu\}$, each element of which represents a possible way in which the agent may choose to revise its beliefs. If we take a similar non-deterministic approach in the definition of foundational operators, therefore, a singleton $\Sigma_\psi \downarrow \neg\mu$ coincides with the notion of *deterministic* foundational revision. The connection established by theorem 6 between the latter and postulate (R6) is interesting because it is this postulate that distinguishes the AGM operators from the more general class of coherence operators. Note that, by theorems 2 and 6, deterministic foundational revision satisfies the AGM postulates, and thus, by theorem 5, it can be captured with a TO foundations operator. And since TO operators are easily seen to be

³Had we kept the syntax-based definition of foundational revision, only theorems 2 and 3 would have to be marginally weakened, since syntax-based operators satisfy only a weaker form of (R4). Theorems 4 and 5 would remain unaffected, as choosing a set of basic beliefs is the same as choosing an axiomatization.

deterministic in the sense of theorem 6, it follows that the TO condition (equivalently, unambiguous prioritized revision) completely characterizes deterministic foundational operators.⁴

Some implications for non-monotonic reasoning

It is well known (Gärdenfors 1991; Katsuno & Satoh 1991; Arlo-Costa & Shapiro 1992) that there is a close connection between belief revision and certain frameworks for non-monotonic reasoning. The natural correlate of coherence revision operators is the preferential logic of (Shoham 1987), while the natural correlate of foundational revision operators is Poole’s (1988) system of default reasoning. As we now show, it is an easy consequence of our results that both non-monotonic frameworks are equally expressive, in the sense that they can capture exactly the same set of non-monotonic consequence relations.

Recall that a (propositional) *preferential consequence relation* is a relation \vdash_{\leq} defined by $\mu \vdash_{\leq} \theta$ iff $Min(Mod(\mu), \leq) \subseteq Mod(\theta)$, where \leq is a (possibly total) preorder over the set \mathcal{W} of interpretations of the propositional language \mathcal{L} . It is easy to see from theorem 1 that a coherence revision operator \circ induces a preferential relation \vdash_{\leq} for every formula ψ , satisfying

$$\mu \vdash_{\leq} \theta \text{ iff } \theta \in Cn(\psi \circ \mu). \quad (1)$$

(Set $\leq = \leq_\psi$ in \vdash_{\leq} , where \leq_ψ is the preorder associated to ψ by the operator \circ .) And conversely, given a preferential relation \vdash_{\leq} , there exists a formula ψ and a coherence revision operator \circ satisfying expression 1 (choose ψ so that $Mod(\psi) = Min(\mathcal{W}, \leq)$, and choose \circ so that $\leq_\psi = \leq$).

Similarly, there is a very close connection between Poole’s “default theories” and foundational belief revision. Recall that a default theory in the sense of Poole is a pair $\langle D, F \rangle$, where $D, F \subseteq_{fin} \mathcal{L}$ are, respectively, a set of “defaults” and a consistent set of “facts.” The (cautious) non-monotonic consequence relation defined by Poole, written \vdash_{\sim} , is defined by $\langle D, F \rangle \vdash_{\sim} \theta$ iff $\theta \in \bigcap_{\Gamma \in D \downarrow \neg \bigwedge F} Cn(\Gamma \cup F)$, or, equivalently, as shown in (Nebel 1989), iff $\theta \in Cn(D \circ_S \bigwedge F)$, where \circ_S is the basic syntax based revision operator.

This is then an easy consequence of our results:

Theorem 7 *For any preferential consequence relation \vdash_{\leq} there exists $D \subseteq_{fin} \mathcal{L}$ such that for every $\mu, \theta \in \mathcal{L}$, $\mu \vdash_{\leq} \theta$ iff $\langle D, \{\mu\} \rangle \vdash_{\sim} \theta$. And conversely, for any $D \subseteq_{fin} \mathcal{L}$ there exists a preferential relation \vdash_{\leq} such that for any $F \subseteq_{fin} \mathcal{L}$, $\theta \in \mathcal{L}$, $\bigwedge F \vdash_{\leq} \theta$ iff $\langle D, F \rangle \vdash_{\sim} \theta$.*

⁴To verify the equivalence of TO and unambiguous prioritized revision, note: any arbitrary unambiguous prioritized ordering can be imposed on the set of basic beliefs of a TO operator without affecting the result of revision; conversely, because unambiguous prioritized revision is deterministic, it can be captured by a TO operator.

Note that this theorem has as a special case “rational” consequence relations, preferential relations \vdash_{\leq} in which \leq is a total preorder. Note also that the “preferred subtheories” framework proposed in (Brewka 1989) does not extend the expressivity of Poole’s framework, as it can be mapped in the same way to a preferential consequence relation that, by the previous theorem, can be captured in Poole’s framework.

Discussion

We have shown that the coherence and foundational theories of belief revision are mathematically equivalent. More precisely, for a finitary propositional language, the family of coherence revision operators defined in the text, which include the AGM operators, and a slightly modified version of syntax-based revision, are equivalent. This modification of the latter is formally trivial, but in our view is well-motivated by some drawbacks of the syntax-based approach, and captures the essence of the foundational theory better than the latter. We have also shown that preferential logic and Poole’s default theories are expressively equivalent, in the sense that they can capture exactly the same non monotonic consequence relations.

In (del Val & Shoham 1994), we encode belief revision in a situation calculus enriched with epistemic operators and with a knowledge-gathering action for learning new information. In this framework, which inspired most of the results of this paper, update and revision can be jointly captured in a way that makes the temporal evolution of the agent’s beliefs explicit, and both styles of revision, as well as the associated forms of non-monotonic reasoning, can be captured, all within a circumscriptive framework for reasoning about action.

References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic* 50:510–530.
- Arlo-Costa, H. L., and Shapiro, S. J. 1992. Maps between non-monotonic and conditional logic. In *Proc. Third Int. Conf. on Principles of Knowledge Representation and Reasoning*.
- Benferhat, S.; Cayrol, C.; Dubois, D.; Lang, J.; and Prade, H. 1993. Inconsistency management and prioritized syntax-based entailment. In *Proc. Thirteenth Int. Joint Conf. on Artificial Intelligence*.
- Brewka, G. 1989. Preferred subtheories: An extended logical framework for default reasoning. In *Proc. Eleventh Int. Joint Conf. on Artificial Intelligence*.
- Dalal, M. 1988. Investigations into a theory of knowledge base revision. In *Proc. Seventh Conf. of the AAAI*.
- del Val, A., and Shoham, Y. 1994. A unified view of belief revision and update. *Journal of Logic and Computation*. Special Issue on Actions and Processes, M. Georgeff (ed.), to appear.
- Doyle, J. 1979. A truth maintenance system. *Artificial Intelligence* 12:231–272.
- Doyle, J. 1991. Rational belief revision (preliminary report). In *Proc. Second Int. Conf. on Principles of Knowledge Representation and Reasoning*.
- Doyle, J. 1992. Reason maintenance and belief revision: Foundations vs. coherence theories. In Gärdenfors, P., ed., *Belief Revision*. Cambridge University Press.
- Elkan, C. 1990. A rational reconstruction of nonmonotonic truth maintenance systems. *Artificial Intelligence* 43:219–234.
- Fagin, R.; Kuper, G. M.; Ullman, J. D.; and Vardi, M. Y. 1986. Updating logical databases. *Advances in Computing Research* 3.
- Fagin, R.; Ullman, J. D.; and Vardi, M. Y. 1983. On the semantics of updates in databases. In *Proc. Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*.
- Gärdenfors, P. 1988. *Knowledge in Flux*. The MIT Press.
- Gärdenfors, P. 1990. The dynamics of belief systems: Foundations vs. coherence theories. *Revue Internationale de Philosophie* 172:24–46.
- Gärdenfors, P. 1991. Non monotonic reasoning based on expectations. In *Proc. Second Int. Conf. on Principles of Knowledge Representation and Reasoning*.
- Ginsberg, M. L. 1986. Counterfactuals. *Artificial Intelligence* 30:35–79.
- Grosz, B. 1991. Generalizing prioritization. In *Proc. Second Int. Conf. on Principles of Knowledge Representation and Reasoning*.
- Grove, A. 1988. Two modelings for theory change. *Journal of Philosophical Logic* 17:157–170.
- Katsuno, H., and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52:263–294.
- Katsuno, H., and Satoh, K. 1991. A unified view of consequence relations, belief revision and conditional logic. In *Proc. Twelfth Int. Joint Conf. on Artificial Intelligence*.
- Makinson, D. 1985. How to give it up: A survey of some formal aspects of the logic of theory change. *Synthese*.
- Nebel, B. 1989. A knowledge level analysis of belief revision. In *Proc. First Int. Conf. on Principles of Knowledge Representation and Reasoning*.
- Nebel, B. 1991. Belief revision and default reasoning: Syntax-based approaches. In *Proc. Second Int. Conf. on Principles of Knowledge Representation and Reasoning*.
- Pimentel, S. G., and Rodi, W. L. 1991. Belief revision and paraconsistency in a logic programming framework. In *Proc. First Int. Workshop on Logic Programming and Non Monotonic Reasoning*.
- Poole, D. 1988. A logical framework for default reasoning. *Artificial Intelligence* 36:27–47.
- Reinfrank, M.; Dressler, O.; and Brewka, G. 1989. On the relation between truth maintenance and autoepistemic logic. In *Proc. Eleventh Int. Joint Conf. on Artificial Intelligence*.
- Shoham, Y. 1987. A semantical approach to non-monotonic logics. In Ginsberg, M. L., ed., *Readings in Non-Monotonic Reasoning*. Morgan Kaufmann.