

Is it Impossible to Keep Up to Date?[†]

M. Morreau

IMS, Universität Stuttgart, Keplerstr 17

D-7000 Stuttgart 1

mimo@adler.philosophie.uni-stuttgart.de

H. Rott

Fachgruppe Philosophie, Universität Konstanz,

Postfach 5560, D-7750 Konstanz 1

pirott@dknkurz1.bitnet

Abstract

Two different types of updates of information states are distinguished corresponding to two different kinds of incoming information: information about the actual state of domain; and information about changes made to the state of that domain. Gärdenfors [1986] proves a theorem showing that the first kind of update function can have certain intuitively appealing properties only on pain of triviality. Here a similar trivialisation result is proved for the second kind of update function.

I Learning and Adaption

Belief and knowledge would be easier to maintain if the world were not in a constant state of flux, but it is. Because it is, beliefs cannot just accumulate but must constantly be revised so as to keep them up to date. Consider the following very general kind of situation: some *reasoning agent* is trying to maintain a body of truths about some domain, or “world”. Lets say it is you. Your task is complicated by the fact that in this world, another agent, the *actor*, is busy bringing about all kinds of changes. The actor might be an individual, a computer program, a corporation or a rugby team. Two kinds of information are available to you which can help you to keep up to date. First, you obtain facts — or what you take to be facts — about the actual state of the world at any given time (the radio announcer says “Right now, the All Blacks are up by 3:0 ”). Second, you obtain information about changes in the world which the other agent has brought about (“The All Blacks have just scored a second try”). Both kinds of incoming information will lead you to *update* your information state; we reserve the term *learning* for updating in response to incoming information that some state of affairs ϕ

[†] This paper is an improved and shortened version of an unpublished manuscript by the first author. The second author sharpened up the main result of that paper by simplifying to the version presented here.

holds. This is to distinguish it from updating in response to incoming information that some state of affairs ϕ has been brought about, which will be called *adapting*.

It might be thought, and it has been thought (see e.g. Jeffrey [1965, pp. 73-74]), that learning and adapting amount to the same thing. But as has been argued in Katsuno and Mendelzon [1990]⁰ they are distinct update functions, with distinct formal properties. In order to bring out some of the differences which are of importance here, it is convenient to introduce a simple model for reasoning about change. Bodies of belief or *information states* are to be modelled as *theories* (T, T_1, \dots) , which are sets of sentences closed under a supraclassical monotonic consequence operation C_n . We assume that $C_n\{\}$ is among our stock of information states. Let us say that T is *consistent* just in case for no sentence ϕ : $\{\phi, \sim\phi\} \subseteq C_n(T)$; that ϕ is consistent just in case $\{\phi\}$ is consistent; and that ϕ is *consistent with* T just in case $T \cup \{\phi\}$ is consistent. The sentences (ϕ, ψ, ξ, \dots) are taken from a propositional language L , with the usual connectives \sim , $\&$, and \leftrightarrow .

Learning is to be modelled by a learning function $*$, which assigns to each theory T and sentence ϕ a theory $T*\phi$. This theory stands for the information state you would end up in if, while in the information state T , you were to learn that the world is such that ϕ is true. Learning is the subject of the Bayesian tradition in statistics and philosophy (Jeffrey [1965]). Bayesian information states are *probability functions* over L , and learning is modelled as *conditionalisation*. Our more simple modelling is taken from the broadly Bayesian theory of *belief revision* due to the philosopher Peter Gärdenfors [1988]. Indeed, our theories may be thought of as the *tops* of such probability functions, that is the sets of sentences to which they assign unit probability. Nothing in the following will be misunderstood if $*$ is simply thought of as conditionalisation.

Adaption is to be modelled by means of a function $+$, which assigns to each theory T and sentence ϕ a theory $T+\phi$. This theory stands for the information state you would end up in if, while in the information state T , you were to find out that the world has been changed in such a way as to make ϕ true. That is, the models of $T+\phi$ stand for the ways the world might end up, as far as you can tell from within T , after ϕ is brought about by the agent.

Some of the properties of belief revision functions investigated by Gärdenfors are common to both learning and adaption. One property which update functions may have, for example, is that of monotonicity.

DEFINITION: An update function f is *monotonic* iff for all T_1, T_2 and ϕ :

$$T_1 \subseteq T_2 \quad \Rightarrow \quad f(T_1, \phi) \subseteq f(T_2, \phi).$$

f is called *weakly monotonic* iff the same condition holds for all T_1, T_2 and ϕ consistent with T_2 .

⁰ Katsuno and Mendelzon refer to these notions as revising and updating. We find it hard to remember which is which however, and prefer our more suggestive names.

That learning may reasonably be supposed monotonic as long as ϕ is consistent with T_2 is clear, given Gärdenfors' [1988, Section 3.2] argument that $T^*\phi$ and $C_n(T \cup (\phi))$ are to be equated whenever ϕ is consistent with T . Learning, then, is weakly monotonic. But it does not seem wise to claim full monotonicity for $*$ (compare Gärdenfors [1988, pp. 59-60]). That adaption may reasonably be supposed fully monotonic is argued properly in Morreau [1991], but it is also intuitively very plausible: the more you know about what the world is like before some change is made, the more you will know about what it is like after the change is made. We believe that monotonicity for $+$ has quite as good a motivation as weak monotonicity, but we wish to add it to the records that the weak notion suffices to prove the trivialisation result to come.

Other of the properties of belief revision functions investigated by Gärdenfors cannot reasonably be had by both learning and adaption. The main difference between learning and adaption is easily illustrated. Suppose you have before you a basket, which you believe to contain either apples or bananas, but not both. This is all you believe about its contents; in particular you have no belief that it does not contain apples. Suppose you were to learn that in fact there are apples in the basket. Then the appropriate way to update your beliefs to accommodate this new information would be to exclude a possibility that you had previously been reckoning with, namely that there were bananas, not apples, in the basket. This way of revising your beliefs is cumulative; in particular, you retain the belief that there is just one kind of fruit in the basket.

Suppose on the other hand that, instead of learning that in fact there are apples in the basket, you were to see someone come along and put some apples into the basket. How are you to adapt your beliefs this time, to accommodate this new information that there are apples in the basket? You will reason as follows by cases: beforehand there were either apples and no bananas in there, or bananas but no apples. If the former then now, after the addition of more apples, there will still be apples and no bananas. If the latter, then now there will be both apples and bananas. So now, you will say, either there are apples but no bananas in the basket, or there are both fruits. Note that this time you drop your previous belief that there is just one kind of fruit in the basket.

The crucial difference between learning and adaption illustrated by this example is that the former intuitively must, and the latter intuitively must not, have Gärdenfors' property of expansion:

DEFINITION: An update function f is *expansive* iff for all T and ϕ :

$$\neg\phi \in T \Rightarrow T \subseteq f(T, \phi).$$

If $T^*\phi$ and $C_n(T \cup (\phi))$ are identified whenever ϕ is consistent with T , it follows trivially with the monotonicity of C_n that $*$ is expansive.

II Learning that Something is Wrong

A second important difference between learning and adaptation concerns what happens when ϕ is not consistent with T . In the case of adapting, there is nothing problematic about a state of affairs being brought about which contradicts information about the previous state of the world. It is in the nature of change that some things come to be true which were previously untrue; adaptation must reflect this, by enabling things to be believed true which were previously believed untrue. In Bayesian theory of learning, however, $*$ is traditionally not defined in this case. Indeed, it is not obvious that there is anything of a general and normative nature to be said about how one must revise ones beliefs about the world if, while believing that ϕ is true, one obtains information that in fact $\sim\phi$ is true. Clearly in such a situation something has gone wrong, but there are any number of possibilities as to what it might be. Ones previous information might be wrong, ones new information might be wrong, or both. Ones memory might be playing up, and so on. It is not obvious that general updating advice can be given which is not trivial and yet appropriate to each of these situations.

Peter Gärdenfors has however attempted to give just such advice, extending $*$ to the general case of theory revision, and formulating rationality constraints governing such general revision functions. Apart from expansion and weak monotonicity, the constraints he places on general revision functions include the following:

DEFINITION: (a) An update function f is *successful* iff for all T and ϕ : $\phi \in f(T, \phi)$.
 (b) An update function f is *preservative of consistency* iff $f(T, \phi)$ is consistent whenever each of T and ϕ , taken individually, is consistent.

As constraints on $*$ these are uncontroversial where ϕ is consistent with T . This is again clear if $T*\phi$ and $Cn(T \cup \{\phi\})$ are equated in this special case. But it turns out that they are anything but uncontroversial when extended to general revision functions. In his 1986 article, and again in Section 7.4 of *Knowledge in Flux*, Gärdenfors shows that a general revision function $*$ can be fully monotonic, expansive, successful and preservative of consistency only if the language is so inexpressive as to be trivial.¹ The main suspects when it comes to assigning responsibility for this result, according to Gärdenfors, are the constraints of monotonicity and expansion. The others he thinks less dubious, even in the more general setting.

Adaption, we have said, may reasonably be supposed monotonic, but may not reasonably be supposed to be expansive. Noticing that only one of these main suspects in Gärdenfors' trivialisation result is present, Katsuno and Mendelzon [1990, p. 9] seem to think that adaption will not be faced with such theoretical difficulties. It is the purpose of the rest of this paper to argue to the contrary. In section III we review briefly some properties which adaption may reasonably be supposed to have. Then in section IV we consider a further constraint on adaption which many have thought plausible in

¹ In Gärdenfors' work, full monotonicity comes in when a conditional connective governed by the so-called Ramsey test is introduced into the object language (see Gärdenfors [1988, Chapter 7]).

one form or another: independent items of information should in some sense be treated independently when updating. We formalise a constraint of *respect for independent information*, and show in section V that in combination with the other properties, triviality ensues.

III Properties of Adaption Functions

The constraints placed on adaption functions in this section reduce for the most part to assumptions about the agents of change being rational in various ways. The first, however, that + satisfy success, already follows from the intended meaning of $T+\phi$ as that information state which one would end up in on learning that the agent has succeeded in bringing ϕ about. If one learns that ϕ has been brought about, then ϕ may be assumed true. Note that this is not the same as assuming the agent to be omnipotent! The next constraint is more interesting:

DEFINITION: + is *efficient* iff for all T and ϕ : $\phi \in T \Rightarrow T = T+\phi$.

Efficient functions do nothing if this is not necessary in order to be successful. Efficiency is quite uncontroversial in a Bayesian setting: conditionalising on a sentence with unit prior leaves a probability distribution unchanged. It is interesting that, in this context of adaption, efficiency is problematic, and in need of a quite different justification. What efficiency amounts to is the following rationality assumption about how the actor brings about ϕ : if it is already true that ϕ , then in bringing ϕ about nothing will be changed. Of course this constitutes quite an idealisation, since real agents sometimes do go to unnecessary lengths in trying to achieve what is already the case, even bringing about the opposite of what they had imagined.² The important point is that it should at least be possible, however unrealistically, to *assume* agents to be ideal in this sense, and to reason about their actions on that assumption. Such reasoning is what an efficient + formalises.

It turns out that several interesting interactions between + and & follow from the constraints of success, efficiency and monotonicity which we have already placed on adaption functions. In fact, even weak monotonicity will do in order to prove the following

FACTS:

- i. $Cn\{\}+\phi = Cn\{\phi\}$
- ii. $Cn\{\}+\phi \subseteq Cn\{\}+\phi\&\psi$
- iii. $(Cn\{\}+\phi)+\psi \subseteq Cn\{\}+\phi\&\psi$

Proof: By success and closure under logic, the claims are trivial when ϕ is inconsistent or ψ is inconsistent with ϕ . So let ϕ be consistent and ψ be consistent with ϕ .

² Recently in New York, the first author wanted to go to 5th Avenue. Thinking he was at the time on 7th, he set off east. In fact he was already on 5th.

- i. By the weak monotonicity of $+$ and the monotonicity of C_n , and by efficiency, $C_n\{\} + \phi \subseteq C_n\{\phi\} + \phi = C_n\{\phi\}$. By success, monotonicity of C_n , and closure under logic, $C_n\{\phi\} \subseteq C_n\{\} + \phi$.
- ii. By the monotonicity of C_n , $C_n\{\} \subseteq C_n\{\phi \& \psi\}$, so by the weak monotonicity of $+$, $C_n\{\} + \phi \subseteq C_n\{\phi \& \psi\} + \phi$. But now by efficiency, success and i, $C_n\{\phi \& \psi\} + \phi = C_n\{\phi \& \psi\} = C_n\{\} + \phi \& \psi$.
- iii. From ii, with the weak monotonicity of $+$, $(C_n\{\} + \phi) + \psi \subseteq (C_n\{\} + \phi \& \psi) + \psi$. But by success and efficiency, $C_n\{\} + \phi \& \psi + \psi = C_n\{\} + \phi \& \psi$. Q.E.D.

The constraint of consistency preservation mentioned above is the last of the constraints investigated by Gärdenfors which we wish to place on adaption functions. Changing the world from a place where ϕ is true into a place where $\sim\phi$ is true does not have the effect that inconsistencies become true; that adaption functions must satisfy consistency preservation is a direct reflection of this.

Here we have only briefly argued that the constraints of monotonicity, success, efficiency and consistency are reasonable in the case of adaption functions. That this is so is further supported by the fact that it is possible to model information states and adaption functions which answer to them in a convincing way. These models for reasoning about action are useful in gaining a better understanding planning, among other things. For details the reader is referred to Morreau [1991].

IV Respect for Independent Information

We turn now to the notion of respect for independent information. We restrict ourselves to respecting the *logical* independence of items of information, but since the incompatibility argument readily generalizes to cover other evidential relations, this restriction does not reduce its force.

DEFINITION: For $n \geq 1$, ϕ_1, \dots, ϕ_n are said to be *independent* just in case every sentence of the form $[\sim]\phi_1 \& \dots \& [\sim]\phi_n$, where the negation may be present at arbitrarily many places, is consistent.

Note that ϕ and ψ are independent just in case ϕ and $\phi \leftrightarrow \psi$ are. Now as a first pass at the notion of respecting independence we might have said that $+$ respects independence just in case for all T and for all independent ϕ, ψ :

$$\phi, \psi \in T \quad \Rightarrow \quad \psi \in T + \sim\phi.$$

Such a constraint would however have been unreasonably strict. Let for example C_n be classical consequence, let p and q be two independent atomic sentences, and define T to be $C_n\{p, q\}$. Obviously $p \leftrightarrow q \in T$, since $p \leftrightarrow q$ follows classically from p and q . Clearly, too, $p \leftrightarrow q$ and $\sim p$ are independent sentences. Now it is in the intuitive spirit of respecting independence that $T + \sim p$ just be $C_n\{\sim p, q\}$. But such $+$ does not respect independence in the formal sense defined above since otherwise $p \leftrightarrow q \in T + \sim p = C_n\{\sim p, q\}$, which is clearly not the case.

A little thought shows why it would however be unreasonable to require that $p \leftrightarrow q$ be carried over into $T + \sim p$ if $+$ is to respect independence: the only reason why $p \leftrightarrow q$ was there in the first place was

that it follows with logic from the two sentences p and q which together define T . It is unsurprising and desirable that information which tags along behind other information in this way should disappear again just as soon as that other information does. Formulating an entirely satisfying notion of respecting independence is likely to be tricky, but for now it is not necessary to do so. All that is needed is that any such notion be stricter than the following version:

DEFINITION: $+$ respects independence iff for all T and independent $\phi_1, \phi_2, \dots, \phi_n$:

- (i) $\phi_j \in T + \phi_1 + \phi_2 + \dots + \phi_n$ and
 (ii) $\phi_j \in T + \phi_1 + \phi_2 + \dots + \phi_n + \sim\phi_i$ for any $i \neq j$.

Thus provided there are no logical relations between a number of sentences, respect for independence requires that they can be added and deleted again independently of each other.

If $+$ respects independence in addition to being monotonic, successful and efficient, then FACT iii. may be reversed for independent ϕ and ψ :

FACT iv. $Cn\{\} + \phi \& \psi \subseteq (Cn\{\} + \phi) + \psi$

Proof: Assume $+$ in addition respects independence, and that ϕ and ψ are independent. By FACT i, $Cn\{\} + \phi \& \psi = Cn\{\phi \& \psi\}$. Now by independence $\phi \in (Cn\{\} + \phi) + \psi$, so with success and closure under logic: $Cn\{\phi \& \psi\} \subseteq (Cn\{\} + \phi) + \psi$.

Q.E.D.

Why should adaption functions respect independence? Like the constraint of efficiency (but unlike the constraints of success and monotonicity) this new constraint reduces to a "rationality assumption" concerning the actor who is changing the world about. But this time the assumption is a very strong one, and in need of some justification. To put it very briefly, the idea is that in pursuing a series of goals, the actor is assumed to avoid undoing previous achievements wherever possible.

Less cryptically, imagine an actor who is not just changing the world at random, but is doing so in response to a series of commands. That is, he has a series of goals ϕ_1, \dots, ϕ_n which he is required to successively bring about. So letting T stand for the reasoning agent's views on what the world was like beforehand, $T + \phi_1 + \phi_2 + \dots + \phi_n$ is taken to stand for his views about what the world is like after the actor has worked through this series of goals. Suppose in addition that the actor is obedient; he might for example be a robot programmed to dutifully carry out commands. Then it makes sense for the robot to maximize the number of commands which is adhered to. In particular, if it is possible for all of ϕ_1, \dots, ϕ_n to be true at the same time, then after the actor has responded to the commands to bring about the truth of successively ϕ_1, \dots, ϕ_n , all of these sentences should be true.

Maximizing the number of goals reached in this manner will in general bring the robot extra work: suppose for example that he is first told to give each of the children at a party just a single piece of cake. He can comply with this command by giving the first child a piece of chocolate cake, the second child a piece of apple cake, and so on; suppose that is in fact his response to this first command. Then comes another command: the second child doesn't like apples, and insists that the robot bring him a piece of chocolate cake. How now to respond? It won't do for the robot to simply give the second child a piece of chocolate cake, even though this would be a way of complying with the second command taken in isolation. For then the child would have two pieces of cake, and the first command would be violated. On this conception, what the robot must do is in addition to take away the piece of apple cake, thus bringing about a state of affairs where both of the commands given to him have been obeyed: each child has just a single piece of cake, and the second child has a piece of chocolate cake.

Respect for independence, as formulated here, constrains the adaption function to assume the actor to be obedient in much this way. For consider any independent series ϕ_1, \dots, ϕ_n of sentences. Then by definition of independence it is possible for *all* of ϕ_1, \dots, ϕ_n to be true at the same time. So after the actor has responded to the commands to bring about the truth of successively ϕ_1, \dots, ϕ_n , all of these sentences should be true. This is what (i) of the above definition of respect for independence requires. As for (ii) of that definition, return for a moment to the children's party. And suppose that the robot has just responded to the second command by giving the second child a piece of chocolate cake, at the same time taking away the apple cake so as to comply with the first command. At that moment comes a third command, this time from the child's mother: he is not to have chocolate cake after all (say because he is allergic to chocolate). In order to comply with this new command the chocolate cake will have to be taken away from the second child again, so obedience compels the robot to do this. But since it is still logically possible for the robot to continue to comply with the first command, maximizing the number of commands complied with compels him to do so. The child must be given some piece of cake (though not chocolate cake) in return, so that it still holds that every child has one piece of cake. Part (ii) of the definition of respect for independence constrains the adaption function to assume the actor to be acting in just this way. Again, the restriction to independent $\phi_1 \dots \phi_n$ is essential, in that only if $\{\phi_1, \phi_2, \dots, \sim\phi_i, \phi_{i+1}, \dots, \phi_n\}$ is a consistent set can (ii) in that definition be expected to hold.

Respect for independent information, understood one way or the other, has held some intuitive appeal for several authors working in the field of belief revision. The above attempt to formalize it, however, and our attempt to justify the formalization, do serve to bring out two points which might go unnoticed in a less formal discussion. First, in introducing the constraint of respect for independent information, we in a way move outside of the simple formal framework of theory revision introduced at the beginning of this paper. For now it is not only the contents of a theory which can decide how it is to be adapted to the changing world, but also its update history. Our interpretation of the trivialisation result to come is that this is what is to be learned from the whole exercise: that if such constraints are to be introduced, then it will be necessary to allow certain statically identical theories to show different dynamic behaviour.

Second, as will be obvious, respect for independence makes quite a strong assumption about the actor. It is not at all difficult to imagine actors bringing about changes without respecting independence. Indeed, the earlier example illustrating that adaption cannot be expected to satisfy expansion serves just as well as an example that adaption cannot in generally be expected to respect independence, since the following two sentences are independent in the required sense: *the basket contains either apples or bananas, but not both*, and *the basket contains apples*³. More evidence of the strength of respect for independence is that it actually implies a good part of expansion:

FACT v. If $+$ respects independence in addition to being monotonic, successful and efficient, then for any *finitely axiomatisable* theory $T = Cn(\phi)$ and any ψ such that $\sim\phi \& \sim\psi$ is consistent: $\sim\psi \notin T \Rightarrow T \subseteq T+\psi$.

Proof: Let $T = Cn(\phi)$, let $\sim\phi \& \sim\psi$ be consistent, and let $\sim\psi \notin T$. We have to show that $T \subseteq T+\psi$. To this end, suppose first that ϕ and ψ are independent. Then we have $T+\psi = Cn(\phi)+\psi =$ (by FACT i) $= Cn(\)+\phi+\psi =$ (by FACTS iii and iv) $Cn(\)+\phi \& \psi =$ (by FACT i) $Cn(\phi \& \psi) =$ (by logic) $Cn(Cn(\phi) \cup \{\psi\}) = Cn(T \cup \{\psi\})$. So $T \subseteq T+\psi$.

Suppose on the other hand that ϕ and ψ are not independent. We distinguish four cases. If ϕ and ψ are inconsistent then $\sim\psi \in T$, contradicting our assumption. If ϕ and $\sim\psi$ are inconsistent then $\psi \in T$, so that by efficiency $T = T+\psi$. If $\sim\phi$ and ψ are inconsistent then $\phi \in Cn(\psi)$, so that $T = Cn(\phi) \subseteq T+\psi$ by succes and closure under logic. Finally, the case where $\sim\phi$ and $\sim\psi$ are inconsistent is excluded by assumption. *Q.E.D.*

The constraint of respect for independence is strong, then, but has at least some plausibility in situations where the more general constraints introduced earlier in this paper would also seem reasonable. The kind of obedience required of the robot in the story above is familiar from artificial intelligence discussions of "reasoning about action" such as Winslett [1988].

V A Trivialisation Result for Adaption

Let us say that L is *non-trivial* just in case there are two independent L sentences ϕ and ψ , and that L is otherwise *trivial*. Then the following can be shown:

FACT v: Suppose that $+$ respects independence, and is successful, efficient, preservative of consistency and weakly monotonic. Then L is trivial.

³ Thanks to our reviewer for pointing this out.

Proof: Let T be the information state $Cn\{\}$, and suppose for reductio that L is non-trivial, in virtue of the two independent sentences ϕ and ψ . Clearly, ϕ and $\phi \leftrightarrow \psi$ too are independent sentences.

Now we have: $T + \phi + \psi =$ (by FACTS iii. and iv.) $T + \phi \& \psi =$ (by FACT i.) $Cn\{\phi \& \psi\} =$ (by logic) $Cn\{\phi \& (\phi \leftrightarrow \psi)\} =$ (by FACT i.) $T + \phi \& (\phi \leftrightarrow \psi) =$ (by FACTS iii. and iv.) $T + \phi + \phi \leftrightarrow \psi$.

It follows that

$$(a) T + \phi + \psi + \sim\phi = T + \phi + \phi \leftrightarrow \psi + \sim\phi.$$

By respect for independent information, then

$$(b) \psi \in T + \phi + \psi + \sim\phi.$$

Similarly, and with (a),

$$(c) \phi \leftrightarrow \psi \in T + \phi + \phi \leftrightarrow \psi + \sim\phi = T + \phi + \psi + \sim\phi.$$

Finally, by success,

$$(d) \sim\phi \in T + \phi + \psi + \sim\phi.$$

Now by (b), (c) and (d), $T + \phi + \psi + \sim\phi$ is an inconsistent set. But this is in contradiction with the assumption that $+$ satisfies consistency preservation, three applications of which tell us that $T + \phi + \psi + \sim\phi$ is consistent.

Q.E.D.

We note that the above proof uses only weak monotonicity, whereas Gärdenfors' trivialisation theorem for learning uses full monotonicity. Recalling that other of the constraints placed on adaption functions have been thought reasonable in the case of learning, the above result may be of interest in that area too.

VI Conclusion

So is it impossible to keep up to date or is it not? It clearly is not. Apparently something is wrong with the simple modelling of adaption presented in the above. But what is this result telling us about how to model it better? We think that information states need to be individuated more finely than was done above, thus allowing states which contain exactly the same information about "the world" to have different dynamic behaviour. This could be worked out in several different ways. One would be to model information states not as theories but as the *belief bases*. Another would be to equip theories with relations of *epistemic entrenchment* which are designed to regulate changes of belief. A third way, related to the second one (cf. Rott [1991]), would be to keep with the simple modelling of information states as theories, but to introduce into the object language sentences, say *conditionals* interpreted according to the Ramsey rule, in which the differing dynamic behaviour of extensionally identical states is expressed. The constraints we placed on adaption functions will have to be reconsidered in the extended settings.

VII Literature

- Gärdenfors, P. "Belief Revisions and the Ramsey Test for Conditionals", *Philosophical Review*, 95, 1986.
- Gärdenfors, P. *Knowledge in Flux*, Bradford Books, MIT Press, Cambridge, Mass., 1988.
- Jeffrey, R.C. *The Logic of Decision*, McGraw Hill, New York, 1965.
- Katsuno, H., & Mendelzon, A. "On the Difference Between Updating a Knowledge Database and Revising It", Technical Report, University of Toronto, 1990.
- Morreau, M. "Planning from First Principles", in P. Gärdenfors (ed.), *Belief Revision*, Cambridge University Press, Cambridge 1991.
- Rott, H. "A Nonmonotonic Conditional Logic for Belief Revision I", in A. Fuhrmann and M. Morreau (eds.), *The Logic of Theory Change*, Springer, Berlin, 1991.
- Winslett, M. "Reasoning about action using a possible models approach", in Proceedings of *AAAI*, 1988.