Action Theory Revision in Dynamic Logic

Ivan José Varzinczak

IRIT – Université de Toulouse
Toulouse, France
ivan.varzinczak@irit.fr

Meraka Institute
CSIR, Pretoria, South Africa
ivan.varzinczak@meraka.org.za

Abstract

Like any other logical theory, action theories in reasoning about actions may evolve, and thus need revision methods to adequately accommodate new information about the behavior of actions. Here we give a semantics that complies with minimal change for revising action theories stated in a version of PDL. We give algorithms that are proven correct w.r.t. the semantics for those theories that are modular.

Introduction

In logic-based approaches to reasoning about actions, theories are collections of statements of the form: "if *context*, then *effect* after *every execution* of *action*" (effect laws); and "if *precondition*, then *action executable*" (executability laws). For example, in Propositional Dynamic Logic (PDL) (Harel, Tiuryn, and Kozen 2000), one could have the law $(\neg p_1 \wedge \neg p_2) \rightarrow [a]p_1$, saying that in every context where $\neg p_1 \wedge \neg p_2$ is the case, after every execution of action a we get the effect p_1; and $(p_1 \vee \neg p_2) \rightarrow \langle a \rangle \top$, stating that $p_1 \vee \neg p_2$ is a sufficient condition for a's executability.

These are examples of what we call *action laws*, as they specify the behavior of the actions of a given domain. Besides that we can also have laws mentioning no action at all (static laws). They characterize the underlying structure of the world, i.e., its possible states. For instance, having $p_1 \rightarrow p_2$ as a static law would mean $p_1 \wedge \neg p_2$ is a forbidden state. Action theories will then be collections of laws, each of them seen as a global axiom in PDL.

Well, it may happen that such descriptions have to be revised due e.g. to new incoming information about the behavior of the world. In our example, we may learn that the only valid states are those satisfying $p_1 \wedge p_2$, or that action a has always $\neg p_2$ as outcome in $\neg p_2$-contexts, or even that p_1 is enough to guarantee a's executability. Here we are interested in this kind of theory change.

The contributions of the present work are as follows:

- What is the semantics of revising an action theory $\mathcal{T}$ by a law Φ? How to get minimal change, i.e., how to keep as much knowledge about other laws as possible?
- How to syntactically revise an action theory so that its result corresponds to the intended semantics?

Here we answer these questions.

Logical Preliminaries

Action Theories in Dynamic Logic

Our base formalism is PDL without the $*$ operator. Let $\mathfrak{Act} = \{a_1, a_2, \ldots\}$ be the set of *atomic actions* of a domain. To each a there is associated a modal operator $[a]$. We suppose our multimodal logic is independently axiomatized, i.e., the logic is a fusion and there is no interaction between the modal operators (Kracht and Wolter 1991).

$\mathfrak{Prop} = \{p_1, p_2, \ldots\}$ denotes the set of all *propositional constants* or *atoms*. The set of literals is $\mathfrak{Lit} = \{\ell_1, \ell_2, \ldots\}$, where each ℓ_i is either p or $\neg p$, for some $p \in \mathfrak{Prop}$. In case $\ell = \neg p$, we identify $\neg\ell$ with p. By $|\ell|$ we will denote the atom in literal ℓ.

By $\varphi, \psi, \ldots$ we denote *Boolean formulas*, examples of which are $p_1 \rightarrow p_2$ and $\neg p_1 \oplus p_2$. $\mathfrak{Fml}$ is the set of all Boolean formulas. A propositional valuation v is a *maximally consistent* set of literals. We denote $v \Vdash \varphi$ the fact that v satisfies φ. $val(\varphi)$ is the set of all valuations satisfying φ. $\models_{\mathsf{CPL}}$ denotes the classical consequence relation.

With $IP(\varphi)$ we denote the set of *prime implicants* (Quine 1952) of φ. By π we denote a prime implicant, and $atm(\pi)$ is the set of atoms occurring in π. For given ℓ and π, $\ell \in \pi$ abbreviates 'ℓ is a literal of π'.

We denote complex formulas (with modal operators) by $\Phi, \Psi, \ldots$ $\langle a \rangle$ is the dual operator of $[a]$, ($\langle a \rangle \Phi =_{\text{def}} \neg[a]\neg\Phi$). An example of a complex formula is $(p_1 \wedge (p_2 \vee \neg p_3)) \rightarrow [a](\neg p_1 \vee p_3)$.

A PDL-*model* is a tuple $\mathscr{M} = \langle W, R \rangle$ where W is a set of valuations, and R maps action constants a to accessibility relations $R_a \subseteq W \times W$. Given a model $\mathscr{M}$, $\models^{\mathscr{M}}_{w} p$ (p is true at world w of model $\mathscr{M}$) if $w \Vdash p$; $\models^{\mathscr{M}}_{w} [a]\Phi$ if $\models^{\mathscr{M}}_{w'} \Phi$ for every w' s.t. $(w, w') \in R_a$; truth conditions for the other connectives are as usual. By $\mathcal{M}$ we will denote a set of PDL-models. $\mathscr{M}$ is a model of Φ (noted $\models^{\mathscr{M}} \Phi$) if and only if $\models^{\mathscr{M}}_{w} \Phi$ for all $w \in W$. $\mathscr{M}$ is a model of a set of formulas Σ (noted $\models^{\mathscr{M}} \Sigma$) if and only if $\models^{\mathscr{M}} \Phi$ for every $\Phi \in \Sigma$. Φ is a *consequence of* *the global axioms* Σ in all PDL-models (noted $\Sigma \models_{\mathsf{PDL}} \Phi$) if and only if for every $\mathscr{M}$, if $\models^{\mathscr{M}} \Sigma$, then $\models^{\mathscr{M}} \Phi$.

With PDL we can state laws describing the behavior of actions. Following the tradition in the reasoning about actions community, we here distinguish three types of them.

Static Laws A *static law* is a formula $\varphi \in \mathfrak{Fml}$. It characterizes the possible states of the world. The set of all static laws of a domain is denoted by $\mathcal{S}$.

Effect Laws An *effect law for a* is of the form $\varphi \rightarrow [a]\psi$, where $\varphi, \psi \in \mathfrak{Fml}$. Effect laws relate an action to its effects, which can be conditional. The consequent ψ is the effect which always obtains when *a* is executed in a state where the antecedent φ holds. If *a* is a nondeterministic action, then ψ is typically a disjunction. If ψ is inconsistent we have a special kind of effect law that we call an *inexecutability law*. For example, $(\neg p_1 \wedge p_2) \rightarrow [a]\perp$ says that *a* cannot be executed (there is no *a*-transition) in $\neg p_1 \wedge p_2$-contexts. The set of effect laws of a domain is denoted by $\mathcal{E}$.

Executability Laws An *executability law for a* has the form $\varphi \rightarrow \langle a \rangle \top$, with $\varphi \in \mathfrak{Fml}$. It stipulates the context in which *a* is guaranteed to be executable. (In PDL, the operator $\langle a \rangle$ is used to express executability, $\langle a \rangle \top$ thus reads "*a*'s execution is possible".) The set of all executability laws of a domain is denoted by $\mathcal{X}$.

Action Theories $\mathcal{T} = \mathcal{S} \cup \mathcal{E} \cup \mathcal{X}$ is an *action theory*.

Given an action *a*, $\mathcal{E}_a$ (resp. $\mathcal{X}_a$) will denote the set of only those effect (resp. executability) laws about *a*. For the sake of clarity, we here abstract from the frame and ramification problems, and assume $\mathcal{T}$ contains all frame axioms (cf. (Herzig, Perrussel, and Varzinczak 2006) for a contraction approach within a solution to the frame problem).

Elementary Atoms and Prime Valuations

Given $\varphi \in \mathfrak{Fml}$, $E(\varphi)$ denotes the elementary atoms *actually* occurring in φ. For example, $E(\neg p_1 \wedge (\neg p_1 \vee p_2)) = \{p_1, p_2\}$. An atom *p* is *essential* to φ if and only if $p \in E(\varphi')$ for every φ' such that $\models_{\mathsf{CPL}} \varphi \leftrightarrow \varphi'$. For instance, p_1 is essential to $\neg p_1 \wedge (\neg p_1 \vee p_2)$. $E!(\varphi)$ will denote the essential atoms of φ. (If φ is not contingent, i.e., it is a tautology or a contradiction, then $E!(\varphi) = \emptyset$.)

For $\varphi \in \mathfrak{Fml}$, $\varphi*$ is the set of all $\varphi' \in \mathfrak{Fml}$ such that $\varphi \models_{\mathsf{CPL}} \varphi'$ and $E(\varphi') \subseteq E!(\varphi)$. For instance, $p_1 \vee p_2 \notin p_1*$, as $p_1 \models_{\mathsf{CPL}} p_1 \vee p_2$ but $E(p_1 \vee p_2) \not\subseteq E!(p_1)$. Moreover $E(\varphi*) = E!(\varphi*)$, and whenever $\models_{\mathsf{CPL}} \varphi \leftrightarrow \varphi'$, $E!(\varphi) = E!(\varphi')$ and also $\varphi* = \varphi'*$.

Theorem 1 (Least atom-set theorem (Parikh 1999)) $\models_{\mathsf{CPL}} \varphi \leftrightarrow \bigwedge \varphi*$, *and* $E(\varphi*) \subseteq E(\varphi')$ *for every* φ' *s.t.* $\models_{\mathsf{CPL}} \varphi \leftrightarrow \varphi'$.

Thus for each $\varphi \in \mathfrak{Fml}$ there is a unique least set of elementary atoms such that φ may equivalently be expressed using only atoms from that set.[1]

[1] The dual notion (redundant atoms) is addressed in (Herzig and Rifi 1999), with similar purposes.

Given a valuation *v*, $v' \subseteq v$ is a *subvaluation*. For *W* a set of valuations, a subvaluation v' *satisfies* $\varphi \in \mathfrak{Fml}$ modulo *W* (noted $v' \Vdash_W \varphi$) if and only if $v \Vdash \varphi$ for all $v \in W$ such that $v' \subseteq v$. A subvaluation *v essentially satisfies* φ (modulo *W*), noted $v \Vdash^!_W \varphi$, if and only if $v \Vdash_W \varphi$ and $\{|\ell| : \ell \in v\} \subseteq E!(\varphi)$. If $v \Vdash^!_W \varphi$, we call *v* an *essential subvaluation* of φ (modulo *W*).

Definition 1 *Let* $\varphi \in \mathfrak{Fml}$ *and W be a set of valuations. v is a* prime subvaluation *of* φ *(modulo W) if and only if* $v \Vdash^!_W \varphi$ *and there is no* $v' \subseteq v$ *s.t.* $v' \Vdash^!_W \varphi$.

Prime subvaluations of a formula φ are the weakest states of truth in which φ is true. They are just another way of seeing prime implicants of φ. By $base(\varphi, W)$ we denote the set of all prime subvaluations of φ modulo *W*.

Theorem 2 *Let* $\varphi \in \mathfrak{Fml}$ *and W be a set of valuations. Then for all* $w \in W$, $w \Vdash \varphi$ *if and only if* $w \Vdash \bigvee_{v \in base(\varphi, W)} \bigwedge_{\ell \in v} \ell$.

Closeness Between Models

When revising a model, we will perform a change in its structure. Because there can be several different ways of modifying a model (not all of them minimal), we need a notion of distance between models to identify those that are closest to the original one.

As we are going to see in more depth in the sequel, changing a model amounts to modifying its possible worlds or its accessibility relation. Hence, the distance between two PDL-models will depend upon the distance between their sets of worlds and accessibility relations. These here will be based on the *symmetric difference* between sets, defined as $X \dot{-} Y = (X \setminus Y) \cup (Y \setminus X)$.

Definition 2 *Let* $\mathscr{M} = \langle W, R \rangle$ *be a model.* $\mathscr{M}' = \langle W', R' \rangle$ *is* as close to $\mathscr{M}$ *as* $\mathscr{M}'' = \langle W'', R'' \rangle$, *noted* $\mathscr{M}' \preceq_{\mathscr{M}} \mathscr{M}''$, *if and only if*

- *either* $W \dot{-} W' \subseteq W \dot{-} W''$
- *or* $W \dot{-} W' = W \dot{-} W''$ *and* $R \dot{-} R' \subseteq R \dot{-} R''$

(Notice that other distance notions are also possible, like e.g. considering the *cardinality* of symmetric differences.)

Semantics of Revision

Contrary to action theory contraction (Varzinczak 2008a), where we want the negation of some law to become *satisfiable*, in revision we want to make a new law *valid*. This means that one has to eliminate all cases satisfying its negation. This depicts the duality between revision and contraction: whereas in the latter one invalidates a formula by making its negation satisfiable, in the former one makes a formula valid by forcing its negation to be unsatisfiable prior to adding the new law to the theory.

The idea behind our semantics is as follows: we initially have a set of models $\mathcal{M}$ in which a given formula Φ is (potentially) not valid, i.e., Φ is (possibly) not true in every model in $\mathcal{M}$. In the result we want to have only models of Φ. Adding Φ-models to $\mathcal{M}$ is of no help. Moreover, adding models makes us to lose laws: the corresponding resulting theory would be more liberal.

One solution amounts to deleting from $\mathcal{M}$ those models that are not Φ-models. Of course removing only some of them does not solve the problem, we must delete every such a model. By doing that, all resulting models will be models of Φ. (This corresponds to *theory expansion*, when the resulting theory is satisfiable.) However, if $\mathcal{M}$ contains no model of Φ, we will end up with $\emptyset$. Consequence: the resulting theory is inconsistent. (This is the main revision problem.) In this case the solution is to *substitute* each model $\mathscr{M}$ in $\mathcal{M}$ by its *nearest modification* $\mathscr{M}^{\star}_{\Phi}$ that makes Φ true. This lets us to keep as close as possible to the original models we had. But, what if for one model in $\mathcal{M}$ there are several minimal (incomparable) modifications of it validating Φ? In that case we shall consider all of them. The result will also be a *list of models* $\mathcal{M}^{\star}_{\Phi}$, all being models of Φ.

Before defining revision of sets of models, we present what modifications of (individual) models are.

Revising a Model by a Static Law

Consider the model depicted in Figure 1, and suppose we want to revise it by the Boolean formula $p_1 \vee p_2$, i.e., we want such a formula to be a static law.

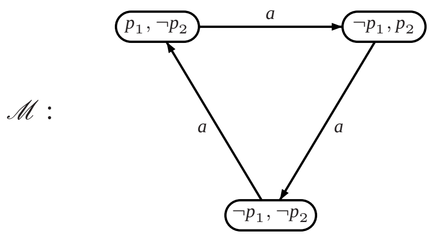

Figure 1: A model where $\neg p_1 \wedge \neg p_2$ is satisfiable.

In such a model, we do not want the formula $\neg p_1 \wedge \neg p_2$ to be satisfiable, so the first step is to remove all worlds in which it is true. The second step is to guarantee that all the remaining worlds satisfy the new law. Such an issue has been largely addressed in the literature on propositional belief base revision and update (Gärdenfors 1988; Winslett 1988; Katsuno and Mendelzon 1992; Herzig and Rifi 1999). Here we can achieve that with a semantics similar to that of classical operators: basically one shall change the set of possible valuations, by removing or adding worlds.

The delicate point in removing worlds is that we may lose some executability laws: in the example, removing $\{\neg p_1, \neg p_2\}$ also removes $p_2 \rightarrow \langle a \rangle \top$. From a semantic point of view, this is intuitive: if the state of the world to which we could move is no longer possible, then we do not have a transition to that state anymore. Hence, if that transition was the only one we had, it is natural to lose it.

Similarly, one could ask what to do with the accessibility relation if new worlds are added (when expansion is not possible): shall new arrows leave/arrive at the new world? If no arrow leaves the new added world, we may lose an executability law. If some arrow leaves it, we may lose an effect law, the same holding if we add an arrow pointing to the new world. If no arrow arrives at this new world, what about the intuition? Do we want to have an unreachable state?

All this discussion shows how drastic a change in the static laws may be: it is a change in the underlying structure (possible states) of the world! Changing it may have as consequence the loss of an effect law or an executability law.

The tradition in the reasoning about actions community says that executability laws are, in general, more difficult to state than effect laws, and hence are more likely to be incorrect. By adding no arrow to the resulting model we here comply with that and postpone correction of executability laws, if needed (cf. (Herzig, Perrussel, and Varzinczak 2006; Varzinczak 2008a)).

The semantics for revision of one model by a static law is as follows:

Definition 3 *Let* $\mathscr{M} = \langle W, R \rangle$. $\mathscr{M}' = \langle W', R' \rangle \in \mathscr{M}^{\star}_{\varphi}$ *if and only if:*

- $W' = (W \setminus val(\neg\varphi)) \cup val(\varphi)$
- $R' \subseteq R$

Clearly $\models^{\mathscr{M}'} \varphi$ for each $\mathscr{M}' \in \mathscr{M}^{\star}_{\varphi}$. The minimal models resulting from revising a model $\mathscr{M}$ by φ are those closest to $\mathscr{M}$ w.r.t. $\preceq_{\mathscr{M}}$:

Definition 4 $revise(\mathscr{M}, \varphi) = \bigcup \min\{\mathscr{M}^{\star}_{\varphi}, \preceq_{\mathscr{M}}\}$

Revising a Model by an Effect Law

Let our language now have three atoms and consider the model $\mathscr{M}$ in Figure 2.

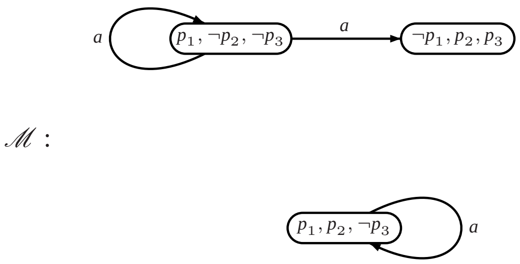

Figure 2: A model where $p_1 \wedge \langle a \rangle p_2$ is satisfiable.

(Notice that $\models^{\mathscr{M}} p_2 \rightarrow p_1 \oplus p_3$.) Suppose we want to revise $\mathscr{M}$ by $p_1 \rightarrow [a]\neg p_2$. This means that we should guarantee the formula $p_1 \wedge \langle a \rangle p_2$ is satisfiable in none of its worlds. To do that, we have to look at the worlds satisfying it (if any) and either make p_1 false, or make $\langle a \rangle p_2$ false by removing a-arrows leading to p_2-worlds.

In our example, the worlds $\{p_1, \neg p_2, \neg p_3\}$ and $\{p_1, p_2, \neg p_3\}$ satisfy $p_1 \wedge \langle a \rangle p_2$ and both have to change. Flipping p_1 would do the job but also has as consequence the loss of a static law: we would violate $p_2 \rightarrow p_1 \oplus p_3$. Here we think that changing action laws should not have as side effect a change in the static laws. Given their special status, these should change only if explicitly required (see above). In this case, each world satisfying $p_1 \wedge \langle a \rangle p_2$ has to be changed so that $\langle a \rangle p_2$ is no longer true in it. In our example, we should remove the arrows $(\{p_1, \neg p_2, \neg p_3\}, \{\neg p_1, p_2, p_3\})$ and $(\{p_1, p_2, \neg p_3\}, \{p_1, p_2, \neg p_3\})$.

The semantics of one model revision for the case of a new effect law is:

Definition 5 *Let* $\mathscr{M} = \langle W, R\rangle$. $\mathscr{M}' = \langle W', R'\rangle \in \mathscr{M}^{*}_{\varphi\rightarrow[a]\psi}$ *if and only if:*

- $W' = W$
- $R' \subseteq R$
- *If* $(w, w') \in R \setminus R'$, *then* $\models^{\mathscr{M}}_{w} \varphi$ *and* $\models^{\mathscr{M}}_{w'} \neg\psi$
- $\models^{\mathscr{M}'} \varphi \rightarrow [a]\psi$

The minimal models resulting from the revision of a model $\mathscr{M}$ by a new effect law are those that are closest to $\mathscr{M}$ w.r.t. $\preceq_{\mathscr{M}}$:

Definition 6 *Let* $\mathscr{M}$ *be a model and* $\varphi \rightarrow [a]\psi$ *an effect law. Then* $revise(\mathscr{M}, \varphi \rightarrow [a]\psi) = \bigcup \min\{\mathscr{M}^{*}_{\varphi\rightarrow[a]\psi}, \preceq_{\mathscr{M}}\}$.

Revising a Model by an Executability Law

Let the model depicted in Figure 3 and suppose we want to revise it by the new executability law $p_1 \rightarrow \langle a\rangle\top$.

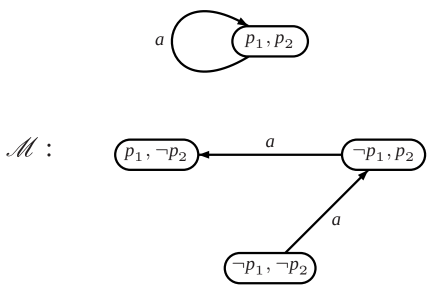

Figure 3: A model where $p_1 \wedge [a]\bot$ is satisfiable.

Observe that $\neg(p_1 \rightarrow \langle a\rangle\top)$ is satisfiable in $\mathscr{M}$, hence we must throw $p_1 \wedge [a]\bot$ away to ensure the new formula is true. To remove $p_1 \wedge [a]\bot$ we have to look at all worlds satisfying it and modify $\mathscr{M}$ so that they no longer satisfy the formula. Given world $\{p_1, \neg p_2\}$, we have two options: change the interpretation of p_1 or add a new arrow leaving this world. A question that raises is 'what choice is more drastic: change a world or an arrow'? Again, here we think that changing the world's content (the valuation) is more drastic, as the existence of such a world was foreseen by some static law and is hence assumed to be as it is, unless we have information supporting the contrary (see above). Thus we shall add a new *a*-arrow from $\{p_1, \neg p_2\}$. Having agreed on that, the issue now is: to which world should the new arrow point? Four options show up: point the arrow to $\{p_1, p_2\}$, $\{\neg p_1, p_2\}$, $\{\neg p_1, \neg p_2\}$ or $\{p_1, \neg p_2\}$ itself. The resulting model is such that the unwanted formula is unsatisfiable and $p_1 \rightarrow \langle a\rangle\top$ holds in all its worlds.

Whereas all these options make the new law true in the resulting model, not all of them comply with minimal change. To witness, putting an *a*-arrow from $\{p_1, \neg p_2\}$ to $\{\neg p_1, \neg p_2\}$ or $\{p_1, \neg p_2\}$ makes us lose the effect law $\neg p_2 \rightarrow [a]p_2$; and pointing it to $\{\neg p_1, p_2\}$ also deletes from the model $p_1 \rightarrow [a]p_1$. Note that these laws are preserved if we point the arrow to $\{p_1, p_2\}$. What would support the choice for the latter?

When pointing a new arrow leaving a world w we want to preserve as many effects as we had before doing so. To achieve this, it is enough to preserve old effects only in w (because the remaining structure of the model remains unchanged after adding *this* new arrow). The operation we must carry out is to observe what is true in w and in the candidate target world w':

- What changes from w to w' ($w' \setminus w$) must be what is obliged to do so.
- What does not change from w to w' ($w \cap w'$) must be what is either obliged or allowed to do so.

This means that every change outside what is forced to change is not an intended one. In our example, when putting the *a*-arrow from $\{p_1, \neg p_2\}$ to $\{\neg p_1, p_2\}$, $\neg p_1$ becomes a possible effect of *a*. As far as $\neg p_1$ is never caused by *a*, there is no justification for having it in a target world of $\{p_1, \neg p_2\}$. Similarly, we want the literals preserved in the target world to be *at most* those that either are consequences of some effect or are usually preserved in that context. Every preservation outside those may make us lose some law. For instance, when putting the new *a*-arrow from $\{p_1, \neg p_2\}$ to $\{\neg p_1, \neg p_2\}$, $\neg p_2$ is preserved. Because $\neg p_2$ is not a necessary effect of *a* and is moreover never preserved across *a*'s execution (in $\mathscr{M}$), there is no reason to preserve it in this new *a*-transition.

This looks like prime implicants, and that is where prime subvaluations play their role: the worlds to which the new arrow shall point are those whose difference w.r.t. the departing world are literals that are relevant, and whose similarity w.r.t. it are literals that we know do not change.

Before giving a formal definition for that, we need to consider two important issues: First, when checking satisfaction of these two conditions, looking just at what is true in the model $\mathscr{M}$ we want to modify is not enough. It can be a model in which a contingent, i.e., not true in all models formula is true. Hence we shall consider all the models in $\mathcal{M}$. Second, if *a* is never executable in w, i.e., $R_a(w) = \emptyset$ for every $\mathscr{M} = \langle W, R\rangle \in \mathcal{M}$, then lots of effects for *a* trivially hold in w, and then not all of them should be taken into account in deciding what has to be changed or preserved. In this case, one should instead look at the effects that hold for those worlds w such that $R_a(w) \neq \emptyset$ (because everything that holds in these worlds also holds trivially in those worlds with no transition by *a*).

Definition 7 *Let* $\mathscr{M} = \langle W, R\rangle$ *be a model,* $w, w' \in W$, $\mathcal{M}$ *a set of models such that* $\mathscr{M} \in \mathcal{M}$, *and* $\varphi \rightarrow \langle a\rangle\top$ *an executability law. Then* w' *is a* relevant target world *of* w *w.r.t.* $\varphi \rightarrow \langle a\rangle\top$ *for* $\mathscr{M}$ *in* $\mathcal{M}$ *if and only if:*

- $\models^{\mathscr{M}}_{w} \varphi$
- *If there is* $\mathscr{M}' = \langle W', R'\rangle \in \mathcal{M}$ *such that* $R'_a(w) \neq \emptyset$*:*
  - *for all* $\ell \in w' \setminus w$, *there is* $\psi' \in \mathfrak{Fml}$ *s.t. there is* $v' \in base(\psi', W)$ *s.t.* $v' \subseteq w'$, $\ell \in v'$, *and for every* $\mathscr{M}_i \in \mathcal{M}$, $\models^{\mathscr{M}_i}_{w} [a]\psi'$
  - *for all* $\ell \in w \cap w'$, *either there is* $\psi' \in \mathfrak{Fml}$ *s.t. there is* $v' \in base(\psi', W)$ *s.t.* $v' \subseteq w'$, $\ell \in v'$, *and for all* $\mathscr{M}_i \in \mathcal{M}$, $\models^{\mathscr{M}_i}_{w} [a]\psi'$*; or there is* $\mathscr{M}_i \in \mathcal{M}$ *s.t.* $\not\models^{\mathscr{M}_i}_{w} [a]\neg\ell$

- *If* $R'_a(w) = \emptyset$ *for every* $\mathscr{M}' = \langle W', R' \rangle \in \mathcal{M}$*:*
  - *for all* $\ell \in w' \setminus w$*, there is* $\mathscr{M}_i = \langle W_i, R_i \rangle \in \mathcal{M}$ *s.t. there is* $u, v \in W_i$ *s.t.* $(u, v) \in R_{i_a}$ *and* $\ell \in v \setminus u$
  - *for all* $\ell \in w \cap w'$*, there is* $\mathscr{M}_i = \langle W_i, R_i \rangle \in \mathcal{M}$ *s.t. there is* $u, v \in W_i$ *s.t.* $(u, v) \in R_{i_a}$ *and* $\ell \in u \cap v$*, or for all* $\mathscr{M}_i = \langle W_i, R_i \rangle \in \mathcal{M}$*, if* $(u, v) \in R_{i_a}$*, then* $\neg\ell \notin v \setminus u$

By $RelTgt(w, \varphi \to \langle a \rangle \top, \mathscr{M}, \mathcal{M})$ *we denote the set of all relevant target worlds of* w *w.r.t.* $\varphi \to \langle a \rangle \top$ *for* $\mathscr{M}$ *in* $\mathcal{M}$*.*

The semantics of one model revision by a new executability law is given by:

Definition 8 *Let* $\mathscr{M} = \langle W, R \rangle$*.* $\mathscr{M}' = \langle W', R' \rangle \in \mathscr{M}^*_{\varphi \to \langle a \rangle \top}$ *if and only if:*

- $W' = W$
- $R \subseteq R'$
- *If* $(w, w') \in R' \setminus R$*, then* $w' \in RelTgt(w, \varphi \to \langle a \rangle \top, \mathscr{M}, \mathcal{M})$
- $\models^{\mathscr{M}'} \varphi \to \langle a \rangle \top$

The minimal models resulting from revising a model $\mathscr{M}$ by a new executability law are those closest to $\mathscr{M}$ w.r.t. $\preceq_{\mathscr{M}}$:

Definition 9 *Let* $\mathscr{M}$ *be a model and* $\varphi \to \langle a \rangle \top$ *be an executability law. Then* $revise(\mathscr{M}, \varphi \to \langle a \rangle \top) = \bigcup \min\{\mathscr{M}^*_{\varphi \to \langle a \rangle \top}, \preceq_{\mathscr{M}}\}$*.*

Revising Sets of Models

Now we are ready to define revision of a set of models $\mathcal{M}$ by a new law Φ:

Definition 10 *Let* $\mathcal{M}$ *be a set of models and* Φ *a law. Then*

$$\mathcal{M}^*_{\Phi} = \begin{cases} \mathcal{M} \setminus \{\mathscr{M} : \not\models^{\mathscr{M}} \Phi\}, & \text{if there is } \mathscr{M} \in \mathcal{M} \text{ s.t. } \models^{\mathscr{M}} \Phi \\ \bigcup_{\mathscr{M} \in \mathcal{M}} revise(\mathscr{M}, \Phi), & \text{otherwise} \end{cases}$$

Observe that Definition 10 comprises both *expansion* and *revision*: in the first one, simple addition of the new law gives a satisfiable theory; in the latter a deeper change is needed to get rid of inconsistency.

Syntactic Operators for Revision

We now turn our attention to the syntactical counterpart of revision. Suppose we have an action theory $\mathcal{T}$ and a law Φ we want to revise $\mathcal{T}$ with. If $\mathcal{T} \cup \{\Phi\}$ is satisfiable, adding Φ to $\mathcal{T}$ (expansion) will do the job. Otherwise, if $\mathcal{T} \cup \{\Phi\} \models_{\mathsf{PDL}} \bot$, then we have to modify the laws in $\mathcal{T}$ to accommodate with the new incoming law (proper revision). Our endeavor here is to perform minimal change at the syntactical level. By $\mathcal{T}^*_{\Phi}$ we denote the result of revising $\mathcal{T}$ with Φ.

Revision by a Static Law

Looking at the semantics of revision by Boolean formulas, we see that revising an action theory by a new static law may conflict with the executability laws: some of them may be lost and thus have to be changed as well. The approach here is to preserve as many executabilities as we can in the old possible states. To do that, we look at each possible valuation that is common to the new $\mathcal{S}$ and the old one. Every time an executability used to hold in that state and no inexecutability holds there in the new theory, we make the action executable in such a context. For those contexts not allowed by the old $\mathcal{S}$, we make a inexecutable (cf. the semantics). Algorithm 1 deals with that ($\mathcal{S} \star \varphi$ denotes the classical revision of $\mathcal{S}$ by φ using any standard method from the literature (Winslett 1988; Katsuno and Mendelzon 1992; Herzig and Rifi 1999)).

Algorithm 1 Revision by a static law

input: $\mathcal{T}, \varphi$
output: $\mathcal{T}^*_{\varphi}$
  if $\mathcal{T} \cup \{\varphi\} \not\models_{\mathsf{PDL}} \bot$ **then**
    $\mathcal{T}^*_{\varphi} := \mathcal{T} \cup \{\varphi\}$
  else
    $\mathcal{S}' := \mathcal{S} \star \varphi$, $\mathcal{E}' := \mathcal{E}$, $\mathcal{X}' := \emptyset$
    for all $\pi \in IP(\mathcal{S}')$ **do**
      for all $A \subseteq \overline{atm(\pi)}$ **do**
        $\varphi_A := \bigwedge_{\substack{p_i \in \overline{atm(\pi)} \\ p_i \in A}} p_i \wedge \bigwedge_{\substack{p_i \in \overline{atm(\pi)} \\ p_i \notin A}} \neg p_i$
        if $\mathcal{S}' \not\models_{\mathsf{CPL}} (\pi \wedge \varphi_A) \to \bot$ **then**
          if $\mathcal{S} \not\models_{\mathsf{CPL}} (\pi \wedge \varphi_A) \to \bot$ **then**
            if $\mathcal{T} \models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \to \langle a \rangle \top$ **and** $\mathcal{S}', \mathcal{E}', \mathcal{X} \not\models_{\mathsf{PDL}} \neg(\pi \wedge \varphi_A)$ **then**
              $\mathcal{X}_a' := \{(\varphi_i \wedge \pi \wedge \varphi_A) \to \langle a \rangle \top : \varphi_i \to \langle a \rangle \top \in \mathcal{X}_a\}$
          else
            $\mathcal{E}' := \mathcal{E}' \cup \{(\pi \wedge \varphi_A) \to [a]\bot\}$
    $\mathcal{T}^*_{\varphi} := \mathcal{S}' \cup \mathcal{E}' \cup \mathcal{X}'$

Revision by an Effect Law

When revising a theory by a new effect law $\varphi \to [a]\psi$, we want to eliminate all possible executions of a leading to $\neg\psi$-states. To achieve that, we look at all φ-contexts and every time a transition to some $\neg\psi$-context is not always the case, i.e., $\mathcal{T} \not\models_{\mathsf{PDL}} \varphi \to \langle a \rangle \neg\psi$, we can safely force $[a]\psi$ for that context. On the other hand, if in such a context there is always an execution of a to $\neg\psi$, then we should strengthen the executability laws to make room for the new effect in that context we want to add. Algorithm 2 below does the job.

Revision by an Executability Law

Revising a theory by a new executability law will have as immediate consequence a change in the set of effect laws: all those laws preventing the execution of a shall be weakened. Besides that, in order to comply with minimal change, we shall ensure that in all models of the resulting theory there will be at most *one* transition by a from those worlds in which $\mathcal{T}$ precluded a's execution.

Let $\mathcal{E}_a^{\varphi,\bot}$ denote a minimum subset of $\mathcal{E}_a$ such that $\mathcal{S}, \mathcal{E}_a^{\varphi,\bot} \models_{\mathsf{PDL}} \varphi \to [a]\bot$. In the case the theory is modular (Herzig and Varzinczak 2005) (see further), interpolation guarantees that this set always exists. Moreover, note that there can be more than one such a set, in which case we denote them $(\mathcal{E}_a^{\varphi,\bot})_1, \ldots, (\mathcal{E}_a^{\varphi,\bot})_n$. Let

$$\mathcal{E}_a^- = \bigcup_{1 \leq i \leq n} (\mathcal{E}_a^{\varphi,\bot})_i$$

Algorithm 2 Revision by an effect law

input: $\mathcal{T}, \varphi \rightarrow [a]\psi$
output: $\mathcal{T}^{*}_{\varphi \rightarrow [a]\psi}$
if $\mathcal{T} \cup \{\varphi \rightarrow [a]\psi\} \not\models_{\mathsf{PDL}} \bot$ **then**
$\mathcal{T}^{*}_{\varphi \rightarrow [a]\psi} := \mathcal{T} \cup \{\varphi \rightarrow [a]\psi\}$
else
$\mathcal{T}' := \mathcal{T}$
for all $\pi \in IP(\mathcal{S} \wedge \varphi)$ **do**
for all $A \subseteq \overline{atm(\pi)}$ **do**
$\varphi_A := \bigwedge_{\substack{p_i \in \overline{atm(\pi)} \\ p_i \in A}} p_i \wedge \bigwedge_{\substack{p_i \in \overline{atm(\pi)} \\ p_i \notin A}} \neg p_i$
if $\mathcal{S} \not\models_{\mathsf{CPL}} (\pi \wedge \varphi_A) \rightarrow \bot$ **then**
for all $\pi' \in IP(\mathcal{S} \wedge \neg\psi)$ **do**
if $\mathcal{T}' \models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow \langle a \rangle \pi'$ **then**

$$\mathcal{T}' := \begin{array}{l} (\mathcal{T}' \setminus \mathcal{X}'_a) \cup \\ \{(\varphi_i \wedge \neg(\pi \wedge \varphi_A)) \rightarrow \langle a \rangle \top : \\ \varphi_i \rightarrow \langle a \rangle \top \in \mathcal{X}'_a\} \end{array}$$

$\mathcal{T}' := \mathcal{T}' \cup \{(\pi \wedge \varphi_A) \rightarrow [a]\psi\}$
if $\mathcal{T}' \not\models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow [a]\bot$ **then**
$\mathcal{T}' := \mathcal{T}' \cup \{(\varphi_i \wedge \pi \wedge \varphi_A) \rightarrow \langle a \rangle \top : \varphi_i \rightarrow \langle a \rangle \top \in \mathcal{T}\}$
$\mathcal{T}^{*}_{\varphi \rightarrow [a]\psi} := \mathcal{T}'$

The effect laws in $\mathcal{E}_a^-$ will serve as guidelines to get rid of $[a]\bot$ in each φ-world allowed by the theory: they are the laws to be weakened to allow for $\langle a \rangle \top$.

The idea behind our algorithm is as follows: to force $\varphi \rightarrow \langle a \rangle \top$ to be true in all models of the resulting theory, we visit every possible φ-context allowed by it and make the following operations to ensure $\langle a \rangle \top$ is the case for that context: Given a φ-context, if $\mathcal{T}$ not always precludes a from being executed in it, we can safely force $\langle a \rangle \top$ without modifying other laws. On the other hand, if a is always inexecutable in that context, then we should weaken the laws in $\mathcal{E}_a^-$. The first thing we must do is to preserve all old effects in all other φ-worlds. To achieve that we specialize the above laws to each possible valuation (maximal conjunction of literals) satisfying φ but the actual one. Then, in the current φ-valuation, we must ensure that action a may have any effect, i.e., from this φ-world we can reach any other possible world. We achieve that by weakening the *consequent* of the laws in $\mathcal{E}_a^-$ to the exclusive disjunction of all possible contexts in $\mathcal{T}$. Finally, to get minimal change, we must ensure that all literals in this φ-valuation that are not forced to change are preserved. We do this by stating a conditional frame axiom of the form $(\varphi_k \wedge \ell) \rightarrow [a]\ell$, where φ_k is the above φ-valuation.

Algorithm 3 gives the pseudo-code for that.

Correctness of the Algorithms

Suppose we have two atoms p_1 and p_2, and only one action a. Let the action theory $\mathcal{T}_1 = \{\neg p_2, p_1 \rightarrow [a]p_2, \langle a \rangle \top\}$. The only model of $\mathcal{T}_1$ is $\mathscr{M}$ in Figure 4. Revising such a model by $p_1 \vee p_2$ gives us the models $\mathscr{M}'_i$, $1 \leq i \leq 3$, in Figure 4. Now, revising $\mathcal{T}_1$ by $p_1 \vee p_2$ will give us ${\mathcal{T}_1}^{*}_{p_1 \vee p_2} = \{p_1 \wedge \neg p_2, p_1 \rightarrow [a]p_2\}$. The only model of ${\mathcal{T}_1}^{*}_{p_1 \vee p_2}$ is $\mathscr{M}'_1$

Algorithm 3 Revision by an executability law

input: $\mathcal{T}, \varphi \rightarrow \langle a \rangle \top$
output: $\mathcal{T}^{*}_{\varphi \rightarrow \langle a \rangle \top}$
if $\mathcal{T} \cup \{\varphi \rightarrow \langle a \rangle \top\} \not\models_{\mathsf{PDL}} \bot$ **then**
$\mathcal{T}^{*}_{\varphi \rightarrow \langle a \rangle \top} := \mathcal{T} \cup \{\varphi \rightarrow \langle a \rangle \top\}$
else
$\mathcal{T}' := \mathcal{T}$
for all $\pi \in IP(\mathcal{S} \wedge \varphi)$ **do**
for all $A \subseteq \overline{atm(\pi)}$ **do**
$\varphi_A := \bigwedge_{\substack{p_i \in \overline{atm(\pi)} \\ p_i \in A}} p_i \wedge \bigwedge_{\substack{p_i \in \overline{atm(\pi)} \\ p_i \notin A}} \neg p_i$
if $\mathcal{S} \not\models_{\mathsf{CPL}} (\pi \wedge \varphi_A) \rightarrow \bot$ **then**
if $\mathcal{T}' \models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow [a]\bot$ **then**

$$\mathcal{T}' := \begin{array}{l} (\mathcal{T}' \setminus \mathcal{E}'^{-}_{a}) \cup \\ \{(\varphi_i \wedge \neg(\pi \wedge \varphi_A)) \rightarrow [a]\psi_i : \\ \varphi_i \rightarrow [a]\psi_i \in \mathcal{E}'^{-}_{a}\} \cup \\ \{(\varphi_i \wedge \pi \wedge \varphi_A) \rightarrow [a] \bigoplus_{\substack{\pi' \in IP(\mathcal{S}) \\ A' \subseteq \overline{atm(\pi')}}} (\pi' \wedge \varphi_{A'}) : \\ \varphi_i \rightarrow [a]\psi_i \in \mathcal{E}'^{-}_{a}\} \end{array}$$

for all $L \subseteq \mathfrak{Lit}$ **do**
if $\mathcal{S} \models_{\mathsf{CPL}} (\pi \wedge \varphi_A) \rightarrow \bigwedge_{\ell \in L} \ell$ **then**
for all $\ell \in L$ **do**
if $\mathcal{T} \models_{\mathsf{PDL}} \ell \rightarrow [a]\bot$ **or** ($\mathcal{T} \not\models_{\mathsf{PDL}} \ell \rightarrow [a]\neg\ell$ **and** $\mathcal{T} \models_{\mathsf{PDL}} \ell \rightarrow [a]\ell$) **then**
$\mathcal{T}' := \mathcal{T}' \cup \{(\pi \wedge \varphi_A \wedge \ell) \rightarrow [a]\ell\}$
$\mathcal{T}' := \mathcal{T}' \cup \{(\pi \wedge \varphi_A) \rightarrow \langle a \rangle \top\}$
$\mathcal{T}^{*}_{\varphi \rightarrow \langle a \rangle \top} := \mathcal{T}'$

in Figure 4. This means that the semantic revision produces models (viz. $\mathscr{M}'_2$ and $\mathscr{M}'_3$ in Figure 4) that are not models of the revised theories.

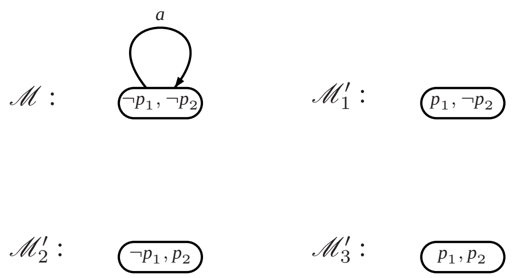

Figure 4: The model $\mathscr{M}$ of $\mathcal{T}$ and the semantic revision of $\mathscr{M}$ by $p_1 \vee p_2$.

The other way round, the algorithms may produce theories whose models do not result from the semantic revision of some model of the original theory. As an example, consider $\mathcal{T}_2 = \{(p_1 \vee p_2) \rightarrow [a]\bot, \langle a \rangle \top\}$, whose only model is $\mathscr{M}$ in Figure 4. The revision of $\mathscr{M}$ by $p_1 \vee p_2$ is as above. However ${\mathcal{T}_2}^{*}_{p_1 \vee p_2} = \{p_1 \vee p_2, (p_1 \vee p_2) \rightarrow [a]\bot\}$ has a model $\mathscr{M}'' = \langle \{\{p_1, p_2\}, \{p_1, \neg p_2\}, \{\neg p_1, p_2\}\}, \emptyset \rangle$ that is not in $\mathscr{M}^{*}_{p_1 \vee p_2}$.

This happens because the possible states are not completely characterized by the static laws in $\mathcal{S}$. Fortunately we get the right result by requiring $\mathcal{S}$ to be 'big enough'. This is connected with the principle of *modularity* (Herzig and Varzinczak 2005):

Definition 11 (Modularity (Herzig and Varzinczak 2005)) $\mathcal{T}$ *is* modular *if and only if for every* $\varphi \in \mathfrak{Fml}$*, if* $\mathcal{T} \models_{\mathsf{PDL}} \varphi$*, then* $\mathcal{S} \models_{\mathsf{CPL}} \varphi$*.*

Under modularity, revision of models of $\mathcal{T}$ by a law Φ in the semantics produces models of the output of the algorithms $\mathcal{T}^{\star}_{\Phi}$:

Theorem 3 *Let* $\mathcal{T}$ *be modular and* Φ *be a law. For all models* $\mathscr{M}'$*, if* $\mathscr{M}' \in \mathcal{M}^{\star}_{\Phi}$*, for some* $\mathcal{M} = \{\mathscr{M} : \models^{\mathscr{M}} \mathcal{T}\}$*, then* $\models^{\mathscr{M}'} \mathcal{T}^{\star}_{\Phi}$*.*

Also under modularity, models of $\mathcal{T}^{\star}_{\Phi}$ result from revision of models of $\mathcal{T}$ by Φ:

Theorem 4 *Let* $\mathcal{T}$ *be modular and* Φ *a law. For every* $\mathscr{M}'$*, if* $\models^{\mathscr{M}'} \mathcal{T}^{\star}_{\Phi}$*, then* $\mathscr{M}' \in \mathcal{M}^{\star}_{\Phi}$*, for some* $\mathcal{M} = \{\mathscr{M} : \models^{\mathscr{M}} \mathcal{T}\}$*.*

In (Herzig and Varzinczak 2005) algorithms are given to check whether $\mathcal{T}$ satisfies the principle of modularity and also to make $\mathcal{T}$ satisfy it, if that is not the case.

Modular theories have other interesting properties (Herzig and Varzinczak 2007): for example, consistency amounts to that of $\mathcal{S}$; deduction of effect laws does not need the executability ones and vice versa; prediction of an effect of a sequence of actions $a_1; \ldots ; a_n$ does not need the effect laws for actions other than $a_1, \ldots, a_n$. This also applies to plan validation when deciding whether $\langle a_1; \ldots ; a_n\rangle\varphi$ is the case.

Conclusion and Perspectives

Contrary to classical belief change, the problem of action theory change has only recently received attention in the literature, both in action languages (Baral and Lobo 1997; Eiter et al. 2005) and in dynamic logic (Herzig, Perrussel, and Varzinczak 2006; Varzinczak 2008a).

Here we have studied what revising action theories by a law means, both in the semantics and at the syntactical level. We have defined a semantics based on distances between models that also captures minimal change w.r.t. the preservation of effects of actions. With our algorithms and the correctness results under modularity we have established the link between the semantics and the syntax, and have also shown that the modularity notion is fruitful. Since modularity is preserved across revision (see Lemma 1 in the appendices), it has to be ensured only once during the evolution of the action theory.

Here we presented the case for revision. In (Varzinczak 2008a) we also define the contraction counterpart of action theory change. There we show that moreover our constructions satisfy all Katsuno and Mendelzon's postulates for contraction (Katsuno and Mendelzon 1992).

Our next step on the subject is to define a general framework in which to revise a theory by *any* formula of the language and not only laws. We believe that such a definition will use as basic operations semantic modifications like those we studied here (addition/removal of arrows and worlds). Hence our constructions will help us in better understanding what revision by a general formula means.

Acknowledgements

The author is thankful to Andreas Herzig and Laurent Perrussel for interesting discussions on the subject of this work.

This work has been partially supported by the government of the FEDERATIVE REPUBLIC OF BRAZIL. Grant: CAPES BEX 1389/01-7.

References

Baral, C., and Lobo, J. 1997. Defeasible specifications in action theories. In *Proc. IJCAI*, 1441–1446.

Eiter, T.; Erdem, E.; Fink, M.; and Senko, J. 2005. Updating action domain descriptions. In *Proc. IJCAI*, 418–423.

Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press.

Harel, D.; Tiuryn, J.; and Kozen, D. 2000. *Dynamic Logic*. MIT Press.

Herzig, A., and Rifi, O. 1999. Propositional belief base update and minimal change. *Artificial Intelligence* 115(1):107–138.

Herzig, A., and Varzinczak, I. 2005. On the modularity of theories. In *Advances in Modal Logic*, volume 5. King's College Publications. 93–109.

Herzig, A., and Varzinczak, I. 2007. Metatheory of actions: beyond consistency. *Artificial Intelligence* 171:951–984.

Herzig, A.; Perrussel, L.; and Varzinczak, I. 2006. Elaborating domain descriptions. In *Proc. ECAI*, 397–401.

Katsuno, H., and Mendelzon, A. 1992. On the difference between updating a knowledge base and revising it. In *Belief revision*. Cambridge. 183–203.

Kracht, M., and Wolter, F. 1991. Properties of independently axiomatizable bimodal logics. *J. of Symbolic Logic* 56(4):1469–1485.

Parikh, R. 1999. Beliefs, belief revision, and splitting languages. In *Logic, Language and Computation*, 266–278.

Quine, W. V. O. 1952. The problem of simplifying truth functions. *American Mathematical Monthly* 59:521–531.

Varzinczak, I. 2008a. Action theory contraction and minimal change. To appear in Proc. KR 2008.

Varzinczak, I. 2008b. Action theory revision. Technical Report IRIT/RT–2008-1–FR, IRIT, Toulouse.

Winslett, M.-A. 1988. Reasoning about action using a possible models approach. In *Proc. AAAI*, 89–93.

Proof of Theorem 3

Let Φ be a law, $\mathscr{M}' \in \mathcal{M}^{\star}_{\Phi}$, and let $\mathcal{T}^{\star}_{\Phi}$ be the output of our algorithms on input theory $\mathcal{T}$ and law Φ.

If $\mathcal{T} \cup \{\Phi\} \not\models_{\mathsf{PDL}} \bot$, then $\mathscr{M}' \in \mathcal{M} \setminus \{\mathscr{M} : \not\models^{\mathscr{M}} \Phi\}$ and $\mathscr{M}'$ is a model of $\mathcal{T}^{\star}_{\Phi} = \mathcal{T} \cup \{\Phi\}$.

Let $\mathcal{T} \cup \{\Phi\} \models_{\mathsf{PDL}} \bot$. We analyze each case.

Let Φ be some $\varphi \in \mathfrak{Fml}$. Then $\mathscr{M}' = \langle W', R'\rangle$ where $W' = (W \setminus val(\neg\varphi)) \cup val(\varphi)$ is minimal w.r.t. W and $R' \subseteq R$ is maximal w.r.t. R, for some $\mathscr{M} = \langle W, R\rangle \in \mathcal{M}$.

As we have assumed the syntactical classical revision operator $\star$ is sound and complete w.r.t. its semantics and is moreover minimal, we have $\models^{\mathscr{M}'} \mathcal{S} \star \varphi$. Because $R' \subseteq R$, $\models^{\mathscr{M}'} \mathcal{E}$. Thus it is enough to show that $\mathscr{M}'$ is a model of the added laws.

Given $(\varphi_i \wedge \pi \wedge \varphi_A) \rightarrow \langle a \rangle \top \in \mathcal{T}^*_\varphi$, for every $w \in W'$, if $\models^{\mathscr{M}'}_w \varphi_i \wedge \pi \wedge \varphi_A$, then $w \in W$ (because $\mathcal{S} \not\models_{\mathsf{CPL}} (\pi \wedge \varphi_A) \rightarrow \bot$). From $w \Vdash \varphi_i$ and $\varphi_i \rightarrow \langle a \rangle \top \in \mathcal{X}_a$, we have $R_a(w) \neq \emptyset$. Suppose $R'_a(w) = \emptyset$. As $\models^{\mathscr{M}'} \mathcal{S} \star \varphi \cup \mathcal{E}$ and R' is maximal, every $\mathscr{M}'' = \langle W'', R'' \rangle$ s.t. $\models^{\mathscr{M}''} \mathcal{S} \star \varphi \cup \mathcal{E}$ is s.t. $R''_a(w) = \emptyset$, and then $\mathcal{S} \star \varphi \cup \mathcal{E} \models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow [a]\bot$. Because $\mathcal{T} \models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow \langle a \rangle \top$, and $\mathcal{S} \not\models_{\mathsf{CPL}} (\pi \wedge \varphi_A) \rightarrow \bot$ and $\mathcal{S} \star \varphi \not\models_{\mathsf{CPL}} (\pi \wedge \varphi_A) \rightarrow \bot$, we get $\mathcal{S} \star \varphi, \mathcal{E}, \mathcal{X} \models_{\mathsf{PDL}} \neg(\pi \wedge \varphi_A)$, and then $(\varphi_i \wedge \pi \wedge \varphi_A) \rightarrow \langle a \rangle \top \notin \mathcal{T}^*_\varphi$. Hence $R'_a(w) \neq \emptyset$, and $\models^{\mathscr{M}'} (\varphi_i \wedge \pi \wedge \varphi_A) \rightarrow \langle a \rangle \top$.

If $(\pi \wedge \varphi_A) \rightarrow [a]\bot \in \mathcal{T}^*_\varphi$, then $\mathcal{S} \models_{\mathsf{CPL}} (\pi \wedge \varphi_A) \rightarrow \bot$. Thus, for every $w \in W'$, if $\models^{\mathscr{M}'}_w \pi \wedge \varphi_A$, $R'_a(w) = \emptyset$ and the result follows.

Let Φ now have the form $\varphi \rightarrow [a]\psi$, for $\varphi, \psi \in \mathfrak{Fml}$. Then $\mathcal{M}' = \langle W', R' \rangle$ for some $\mathscr{M} = \langle W, R \rangle \in \mathcal{M}$ s.t. $W' = W$ and $R' \subseteq R$, where R' is maximal w.r.t. R.

From $W' = W$, $\models^{\mathscr{M}'} \mathcal{S}$. As $R' \subseteq R$, $\models^{\mathscr{M}'} \mathcal{E}$. Because $\mathcal{S} \cup \mathcal{E} \subseteq \mathcal{T}^*_{\varphi \rightarrow [a]\psi}$, it suffices to show that $\mathscr{M}'$ is a model of the added laws.

By definition, $\models^{\mathscr{M}'} \varphi \rightarrow [a]\psi$, and then $\models^{\mathscr{M}'} (\pi \wedge \varphi_A) \rightarrow [a]\psi$ for every $\pi \in IP(\mathcal{S} \wedge \varphi)$.

If $(\varphi_i \wedge \pi \wedge \varphi_A) \rightarrow \langle a \rangle \top \in \mathcal{T}^*_{\varphi \rightarrow [a]\psi}$, then for every $w \in W'$, if $w \Vdash \varphi_i \wedge \pi \wedge \varphi_A$, we have $w \Vdash \varphi_i$. As $w \in W$, and $\varphi_i \rightarrow \langle a \rangle \top \in \mathcal{X}_a$, $R_a(w) = \emptyset$. If $R'_a(w) = \emptyset$, then $w' \Vdash \neg\psi$ for every $w' \in R_a(w)$. Thus as far as we added $(\pi \wedge \varphi_A) \rightarrow [a]\psi$ to $\mathcal{T}^*_{\varphi \rightarrow [a]\psi}$, we must have $\mathcal{T}^*_{\varphi \rightarrow [a]\psi} \models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow [a]\bot$. Hence $R'_a(w) \neq \emptyset$.

Let $(\varphi_i \wedge \bigwedge_{\mathcal{T} \models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow \langle a \rangle \neg\psi} \neg(\pi \wedge \varphi_A)) \rightarrow \langle a \rangle \top \in \mathcal{T}^*_{\varphi \rightarrow [a]\psi}$. For every $w \in W'$, if $\models^{\mathscr{M}'}_w \varphi_i \wedge \bigwedge_{\mathcal{T} \models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow \langle a \rangle \neg\psi} \neg(\pi \wedge \varphi_A)$, then $w \Vdash \varphi_i$, and as $w \in W$ and $\varphi_i \rightarrow \langle a \rangle \top \in \mathcal{X}_a$, we have $R_a(w) \neq \emptyset$. If $R'_a(w) = \emptyset$, because $\models^{\mathscr{M}'} \mathcal{S} \wedge \mathcal{E}$ and R' is maximal, every $\mathscr{M}'' = \langle W'', R'' \rangle$ s.t. $\models^{\mathscr{M}''} \mathcal{S} \wedge \mathcal{E}$ is s.t. $R''_a(w) = \emptyset$. Then $\mathcal{S}, \mathcal{E} \models_{\mathsf{PDL}} \bigwedge_{\ell \in w} \ell \rightarrow [a]\bot$. But then $\mathcal{T} \models_{\mathsf{PDL}} \bigwedge_{\ell \in w} \ell \rightarrow [a]\bot$, and as $\varphi_i \rightarrow \langle a \rangle \top \in \mathcal{X}_a$, $\mathcal{T} \models_{\mathsf{PDL}} \neg(\bigwedge_{\ell \in w} \ell \wedge \varphi_i)$, and then $w \notin W$, a contradiction. Hence $R'_a(w) \neq \emptyset$.

Finally, let Φ be of the form $\varphi \rightarrow \langle a \rangle \top$, for some $\varphi \in \mathfrak{Fml}$. Then $\mathscr{M}' = \langle W', R' \rangle$ for some $\mathscr{M} = \langle W, R \rangle \in \mathcal{M}$ s.t. $W' = W$ and $R' = R \cup R_a^{\varphi,\top}$, with

$$R_a^{\varphi,\top} = \{(w, w') : w' \in RelTgt(w, \varphi \rightarrow \langle a \rangle \top, \mathscr{M}, \mathcal{M})\}$$

such that R' is minimal w.r.t. R.

From $W' = W$, $\models^{\mathscr{M}'} \mathcal{S}$. As $R \subseteq R'$, $\models^{\mathscr{M}'} \mathcal{X}$. As far as $\mathcal{S} \cup \mathcal{X} \subseteq \mathcal{T}^*_{\varphi \rightarrow \langle a \rangle \top}$, it is enough to show that $\mathscr{M}'$ satisfies the added laws.

By definition, $\models^{\mathscr{M}'} \varphi \rightarrow \langle a \rangle \top$, and then $\models^{\mathscr{M}'} (\pi \wedge \varphi_A) \rightarrow \langle a \rangle \top$ for every $\pi \in IP(\mathcal{S} \wedge \varphi)$.

If $(\varphi_i \wedge \pi \wedge \varphi_A) \rightarrow [a](\psi_i \vee \bigoplus_{\substack{\pi' \in IP(\mathcal{S}) \\ A' \subseteq \overline{atm(\pi')}}} (\pi' \wedge \varphi_{A'})) \in \mathcal{T}^*_{\varphi \rightarrow \langle a \rangle \top}$, then for every $w \in W'$, if $w \Vdash \varphi_i \wedge \pi \wedge \varphi_A$, then $w \Vdash \varphi_i$. Because $\models^{\mathscr{M}} \varphi_i \rightarrow [a]\psi_i$, we have $\models^{\mathscr{M}}_{w'} \psi_i$ for all $w' \in W$ s.t. $(w, w') \in R_a$, and then $\models^{\mathscr{M}'}_{w'} \psi_i$ for every $w' \in W'$ s.t. $(w, w') \in R'_a \setminus R_a^{\varphi,\top}$. Now, given $(w, w') \in R_a^{\varphi,\top}$, we have $\models^{\mathscr{M}'}_{w'} \bigoplus_{\substack{\pi' \in IP(\mathcal{S}) \\ A' \subseteq \overline{atm(\pi')}}} (\pi' \wedge \varphi_{A'})$, and the result follows.

Let $(\varphi_i \wedge \bigwedge_{\mathcal{T} \models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow [a]\bot} \neg(\pi \wedge \varphi_A)) \rightarrow [a]\psi_i \in \mathcal{T}^*_{\varphi \rightarrow \langle a \rangle \top}$. For every $w \in W'$, if $\models^{\mathscr{M}'}_w \varphi_i \wedge \bigwedge_{\mathcal{T} \models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow [a]\bot} \neg(\pi \wedge \varphi_A)$, then $w \Vdash \varphi_i$, and as $\models^{\mathscr{M}} \varphi_i \rightarrow [a]\psi_i$, we have $\models^{\mathscr{M}}_{w'} \psi_i$ for all $w' \in W$ s.t. $(w, w') \in R_a$. Thus $\models^{\mathscr{M}'}_{w'} \psi_i$ for every $w' \in W'$ s.t. $(w, w') \in R'_a \setminus R_a^{\varphi,\top}$. Now, if $w \not\Vdash \varphi$, then $R_a^{\varphi,\top} = \emptyset$ and the result follows. Otherwise, if $w \Vdash \varphi$, then $\mathcal{T} \not\models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow [a]\bot$, and then $(\varphi_i \wedge \bigwedge_{\mathcal{T} \models_{\mathsf{PDL}} (\pi \wedge \varphi_A) \rightarrow [a]\bot} \neg(\pi \wedge \varphi_A)) \rightarrow [a]\psi_i$ has not been put in $\mathcal{T}^*_{\varphi \rightarrow \langle a \rangle \top}$, a contradiction.

Let now $(\pi \wedge \varphi_A \wedge \ell) \rightarrow [a]\ell \in \mathcal{T}^*_{\varphi \rightarrow \langle a \rangle \top}$. For every $w \in W'$, if $\models^{\mathscr{M}'}_w \pi \wedge \varphi_A \wedge \ell$, then $\models^{\mathscr{M}'}_w \ell$, and then $\models^{\mathscr{M}}_w \ell$. From $(\pi \wedge \varphi_A \wedge \ell) \rightarrow [a]\ell \in \mathcal{T}^*_{\varphi \rightarrow \langle a \rangle \top}$, we have $\mathcal{T} \models_{\mathsf{PDL}} \ell \rightarrow [a]\bot$ or $\mathcal{T} \not\models_{\mathsf{PDL}} \ell \rightarrow [a]\neg\ell$ and $\mathcal{T} \models_{\mathsf{PDL}} \ell \rightarrow [a]\ell$. In both cases, $\models^{\mathscr{M}}_{w'} \ell$ for every $w' \in R_a(w)$, and then $\models^{\mathscr{M}'}_{w'} \ell$ for every w' s.t. $(w, w') \in R' \setminus R_a^{\varphi,\top}$. It remains to show that $\models^{\mathscr{M}'}_{w'} \ell$ for every $w' \in W'$ s.t. $(w, w') \in R_a^{\varphi,\top}$.

Suppose $\not\models^{\mathscr{M}'}_{w'} \ell$. Then $\neg\ell \in w' \setminus w$. From the construction of $\mathscr{M}'$, there is $\mathscr{M}'' = \langle W'', R'' \rangle \in \mathcal{M}$ s.t. there is $(u, v) \in R''_a$ and $\neg\ell \in v \setminus u$, i.e., $\models^{\mathscr{M}''}_u \ell$ and $\models^{\mathscr{M}''}_v \neg\ell$. From $(u, v) \in R''_a$, we do not have $\mathcal{T} \models_{\mathsf{PDL}} \ell \rightarrow [a]\bot$. From $\models^{\mathscr{M}''}_v \neg\ell$, we do not have $\mathcal{T} \models_{\mathsf{PDL}} \ell \rightarrow [a]\ell$. Thus the algorithm has not put $(\pi \wedge \varphi_A \wedge \ell) \rightarrow [a]\ell$ in $\mathcal{T}^*_{\varphi \rightarrow \langle a \rangle \top}$, a contradiction. ■

Proof of Theorem 4

Lemma 1 *Let Φ be a law. If $\mathcal{T}$ is modular and $\mathcal{T} \cup \{\Phi\} \models_{\mathsf{PDL}} \bot$, then $\mathcal{T}^*_\Phi$ is modular.*

Proof: Let Φ be nonclassical. Suppose $\mathcal{T}^*_\Phi$ is not modular. Then there is $\varphi' \in \mathfrak{Fml}$ s.t. $\mathcal{T}^*_\varphi \models_{\mathsf{PDL}} \varphi'$ and $\mathcal{S}' \not\models_{\mathsf{CPL}} \varphi'$, where $\mathcal{S}'$ is static laws in $\mathcal{T}^*_\Phi$. Suppose $\mathcal{T} \not\models_{\mathsf{PDL}} \varphi'$. Then we must have $\mathcal{T}^*_\Phi \models_{\mathsf{PDL}} \neg\varphi' \rightarrow [a]\bot$ and $\mathcal{T}^*_\Phi \models_{\mathsf{PDL}} \neg\varphi' \rightarrow \langle a \rangle \top$.

Suppose Φ has the form $\varphi \rightarrow [a]\psi$, for $\varphi, \psi \in \mathfrak{Fml}$. Then for all $\varphi \wedge \neg\varphi'$-contexts, as far as $\mathcal{T}^*_\Phi \models_{\mathsf{PDL}} (\varphi \wedge \neg\varphi') \rightarrow [a]\bot$, $(\varphi \wedge \neg\varphi') \rightarrow \langle a \rangle \top \notin \mathcal{T}^*_\Phi$. Then $\mathcal{T}^*_\Phi \models_{\mathsf{PDL}} \varphi'$ if and only if $\mathcal{S}' \models_{\mathsf{CPL}} \varphi'$, a contradiction.

Suppose Φ is of the form $\varphi \rightarrow \langle a \rangle \top$, for $\varphi \in \mathfrak{Fml}$. Then for all $\varphi \wedge \neg\varphi'$-contexts such that $\mathcal{T}^*_\Phi \models_{\mathsf{PDL}} (\varphi \wedge \neg\varphi') \rightarrow \langle a \rangle \top$, $\mathcal{T}^*_\Phi \models_{\mathsf{PDL}} (\varphi \wedge \neg\varphi') \rightarrow [a]\bot$ is impossible as far as $\mathcal{E}^-_a$ has been weakened. Then $\mathcal{T}^*_\Phi \models_{\mathsf{PDL}} \varphi'$ if and only if $\mathcal{S}' \models_{\mathsf{CPL}} \varphi'$, a contradiction.

Hence we have $\mathcal{T} \models_{\mathsf{PDL}} \varphi'$. Because Φ is nonclassical, $\mathcal{S}' = \mathcal{S}$. Then $\mathcal{T} \models_{\mathsf{PDL}} \varphi'$ and $\mathcal{S} \not\models_{\mathsf{CPL}} \varphi'$, and hence $\mathcal{T}$ is not modular.

Let now Φ be some $\varphi \in \mathfrak{Fml}$. Suppose $\mathcal{T}^*_\varphi$ is not modular, i.e., there is $\varphi'' \in \mathfrak{Fml}$ s.t. $\mathcal{T}^*_\varphi \models_{\mathsf{PDL}} \varphi''$ and $\mathcal{S}' = \mathcal{S} \star \varphi \not\models_{\mathsf{CPL}} \varphi''$.

From $\mathcal{S}' \not\models_{\mathsf{CPL}} \varphi''$, there is $v \in val(\mathcal{S}')$ s.t. $v \not\Vdash \varphi''$.

If $v \in val(\mathcal{S})$, as $\mathcal{T}$ is modular, $\mathcal{T} \not\models_{\mathsf{PDL}} \varphi''$. From this and $\mathcal{T}^*_\varphi \models_{\mathsf{PDL}} \varphi''$, we must have $\mathcal{T}^*_\varphi \models_{\mathsf{PDL}} \neg\varphi'' \rightarrow [a]\bot$ and $\mathcal{T}^*_\varphi \models_{\mathsf{PDL}} \neg\varphi'' \rightarrow \langle a\rangle\top$. From the latter, we get $\mathcal{T} \models_{\mathsf{PDL}} \neg\varphi'' \rightarrow \langle a\rangle\top$, and from the first we have $\mathcal{T} \models_{\mathsf{PDL}} \neg\varphi'' \rightarrow [a]\bot$. Putting both results together we get $\mathcal{T} \models_{\mathsf{PDL}} \varphi''$. As $\mathcal{S} \not\models_{\mathsf{CPL}} \varphi''$, we have a contradiction.

If $v \notin val(\mathcal{S})$, then $\mathcal{T}^*_\varphi \not\models_{\mathsf{PDL}} \neg\varphi'' \rightarrow \langle a\rangle\top$, as no executability for context $\neg\varphi''$ has been put into $\mathcal{T}^*_\varphi$. Hence $\mathcal{T}^*_\varphi \not\models_{\mathsf{PDL}} \varphi''$, a contradiction. ∎

Lemma 2 *If $\mathscr{M}_{big} = \langle W_{big}, R_{big}\rangle$ is a model of $\mathcal{T}$, then for every $\mathscr{M} = \langle W, R\rangle$ such that $\models^{\mathscr{M}} \mathcal{T}$ there is a minimal (w.r.t. set inclusion) extension $R' \subseteq R_{big} \setminus R$ such that $\mathscr{M}' = \langle val(\mathcal{S}), R \cup R'\rangle$ is a model of $\mathcal{T}$.*

Proof: See (Varzinczak 2008b). ∎

Lemma 3 *Let $\mathcal{T}$ be modular, and Φ be a law. Then $\mathcal{T} \models_{\mathsf{PDL}} \Phi$ if and only if every $\mathscr{M}' = \langle val(\mathcal{S}), R'\rangle$ such that $\models^{\langle W,R\rangle} \mathcal{T}$ and $R \subseteq R'$ is a model of Φ.*

Proof:
(⇒): Straightforward, as $\mathcal{T} \models_{\mathsf{PDL}} \Phi$ implies $\models^{\mathscr{M}} \Phi$ for every $\mathscr{M}$ such that $\models^{\mathscr{M}} \mathcal{T}$, in particular for those that are extensions of some model of $\mathcal{T}$.

(⇐): Suppose $\mathcal{T} \not\models_{\mathsf{PDL}} \Phi$. Then there is $\mathscr{M} = \langle W, R\rangle$ such that $\models^{\mathscr{M}} \mathcal{T}$ and $\not\models^{\mathscr{M}} \Phi$. As $\mathcal{T}$ is modular, the big model $\mathscr{M}_{big} = \langle W_{big}, R_{big}\rangle$ of $\mathcal{T}$ is a model of $\mathcal{T}$. Then by Lemma 2 there is a minimal extension R' of R w.r.t. R_{big} such that $\mathscr{M}' = \langle val(\mathcal{S}), R \cup R'\rangle$ is a model of $\mathcal{T}$. Because $\not\models^{\mathscr{M}} \Phi$, there is $w \in W$ such that $\not\models^{\mathscr{M}}_w \Phi$. If Φ is some $\varphi \in \mathfrak{Fml}$ or an effect law, any extension $\mathscr{M}'$ of $\mathscr{M}$ is such that $\not\models^{\mathscr{M}'}_w \Phi$. If Φ is of the form $\varphi \rightarrow \langle a\rangle\top$, then $\models^{\mathscr{M}}_w \varphi$ and $R_a(w) = \emptyset$. As any extension of $\mathscr{M}$ is such that $(u, v) \in R'$ if and only if $u \in val(\mathcal{S}) \setminus W$, only worlds other than those in W get a new leaving arrow. Thus $(R \cup R')_a(w) = \emptyset$, and then $\not\models^{\mathscr{M}'}_w \Phi$. ∎

Lemma 4 *Let $\mathcal{T}$ be modular and Φ a law. If $\mathscr{M}' = \langle val(\mathcal{S}'), R'\rangle$ is a model of $\mathcal{T}^*_\Phi$, then there is $\mathcal{M} = \{\mathscr{M} : \models^{\mathscr{M}} \mathcal{T}\}$ s.t. $\mathscr{M}' \in \mathcal{M}^*_\Phi$.*

Proof: Let $\mathscr{M}' = \langle val(\mathcal{S}'), R'\rangle$ be such that $\models^{\mathscr{M}'} \mathcal{T}^*_\Phi$. If $\models^{\mathscr{M}'} \mathcal{T}$, the result follows. Suppose $\not\models^{\mathscr{M}'} \mathcal{T}$. We analyze each case.

Let Φ be of the form $\varphi \rightarrow [a]\psi$, for $\varphi, \psi \in \mathfrak{Fml}$. Let $\mathcal{M} = \{\mathscr{M} : \mathscr{M} = \langle val(\mathcal{S}), R\rangle\}$. As $\mathcal{T}$ is modular, by Lemmas 2 and 3, $\mathcal{M}$ is non-empty and contains only models of $\mathcal{T}$.

Suppose $\mathscr{M}'$ is not a minimal model of $\mathcal{T}^*_{\varphi\rightarrow[a]\psi}$, i.e., there is $\mathscr{M}''$ such that $\mathscr{M}'' \preceq_{\mathscr{M}} \mathscr{M}'$ for some $\mathscr{M} \in \mathcal{M}$. Then $\mathscr{M}'$ and $\mathscr{M}''$ differ only in the effect of a in a given φ-world, viz. a $\pi \wedge \varphi_A$-context, for some $\pi \in IP(\mathcal{S} \wedge \varphi)$ and $\varphi_A = \bigwedge_{\substack{p_i \in \overline{atm(\pi)} \\ p_i \in A}} p_i \wedge \bigwedge_{\substack{p_i \in \overline{atm(\pi)} \\ p_i \notin A}} \neg p_i$ such that $A \subseteq \overline{atm(\pi)}$.

Because $\not\models^{\mathscr{M}'} (\pi \wedge \varphi_A) \rightarrow \langle a\rangle\neg\psi$, we must have $\models^{\mathscr{M}''} (\pi \wedge \varphi_A) \rightarrow \langle a\rangle\neg\psi$, and then $\not\models^{\mathscr{M}''} \varphi \rightarrow [a]\psi$. Hence $\mathscr{M}'$ is minimal w.r.t. $\preceq_{\mathscr{M}}$.

When revising by an effect law, $\mathcal{S}' = \mathcal{S}$. Hence taking the right R and $R_a^{\varphi,\neg\psi}$ such that $\mathscr{M} = \langle val(\mathcal{S}), R\rangle$ and $R' = R \setminus R_a^{\varphi,\neg\psi}$, for some $R_a^{\varphi,\neg\psi} \subseteq \{(w, w') : \models^{\mathscr{M}}_w \varphi, \models^{\mathscr{M}}_{w'} \neg\psi$ and $(w, w') \in R_a\}$, we have $\mathscr{M} \in \mathcal{M}$ and then $\mathscr{M}' \in \mathcal{M}^*_{\varphi\rightarrow[a]\psi}$.

Let Φ have the form $\varphi \rightarrow \langle a\rangle\top$, for $\varphi \in \mathfrak{Fml}$. Let $\mathcal{M} = \{\mathscr{M} : \mathscr{M} = \langle val(\mathcal{S}), R\rangle\}$. As $\mathcal{T}$ is modular, by Lemmas 2 and 3, $\mathcal{M}$ is non-empty and contains only models of $\mathcal{T}$.

Suppose that $\mathscr{M}'$ is not a minimal model of $\mathcal{T}^*_{\varphi\rightarrow\langle a\rangle\top}$, i.e., there is $\mathscr{M}''$ such that $\models^{\mathscr{M}''} \mathcal{T}^*_{\varphi\rightarrow\langle a\rangle\top}$ and $\mathscr{M}'' \preceq_{\mathscr{M}} \mathscr{M}'$ for some $\mathscr{M} \in \mathcal{M}$. Then $\mathscr{M}'$ and $\mathscr{M}''$ differ only on the executability of a in a given φ-world, i.e., a $\pi \wedge \varphi_A$-context, for some $\pi \in IP(\mathcal{S} \wedge \varphi)$ and $\varphi_A = \bigwedge_{\substack{p_i \in \overline{atm(\pi)} \\ p_i \in A}} p_i \wedge \bigwedge_{\substack{p_i \in \overline{atm(\pi)} \\ p_i \notin A}} \neg p_i$, such that $A \subseteq \overline{atm(\pi)}$. This means $\mathscr{M}''$ has no arrow leaving this $\pi \wedge \varphi_A$-world. Then $\models^{\mathscr{M}''} (\pi \wedge \varphi_A) \rightarrow [a]\bot$, and hence $\not\models^{\mathscr{M}''} \varphi \rightarrow \langle a\rangle\top$. Hence $\mathscr{M}'$ is a minimal model of $\mathcal{T}^*_{\varphi\rightarrow\langle a\rangle\top}$ w.r.t. $\preceq_{\mathscr{M}}$.

When revising by executability laws, $\mathcal{S}' = \mathcal{S}$. Thus taking the right R and a minimal $R_a^{\varphi,\top}$ such that $\mathscr{M} = \langle val(\mathcal{S}), R\rangle$ and $R' = R \cup R_a^{\varphi,\top}$, for some $R_a^{\varphi,\top} \subseteq \{(w, w') : \models^{\mathscr{M}}_w \varphi$ and $w' \in RelTgt(w, \varphi \rightarrow \langle a\rangle\top, \mathscr{M}, \mathcal{M})\}$, we get $\mathscr{M} \in \mathcal{M}$ and then $\mathscr{M}' \in \mathcal{M}^*_{\varphi\rightarrow\langle a\rangle\top}$.

Finally, let Φ be some $\varphi \in \mathfrak{Fml}$. Then $\mathscr{M}'$ is such that for every $w \in W'$, if $R'_a(w) \neq \emptyset$, then $w \in val(\mathcal{S})$ and $R_a(w) \neq \emptyset$ for every $\mathscr{M} = \langle W, R\rangle \in \mathcal{M}$. Choosing the right $\mathscr{M} \in \mathcal{M}$ the result follows. ∎

Proof of Theorem 4

Let $\mathcal{T}^*_\Phi$ be the output of our algorithms on input theory $\mathcal{T}$ and law Φ. If $\mathcal{T}^*_\Phi = \mathcal{T} \cup \{\Phi\}$, then $\mathcal{T} \cup \{\Phi\} \not\models_{\mathsf{PDL}} \bot$, and hence every $\mathscr{M}'$ such that $\models^{\mathscr{M}'} \mathcal{T}^*_\Phi$ is such that $\mathscr{M}' \in \mathcal{M} \setminus \{\mathscr{M} : \not\models^{\mathscr{M}} \Phi\}$ and the result follows.

Suppose $\mathcal{T} \cup \{\Phi\} \models_{\mathsf{PDL}} \bot$. From the hypothesis that $\mathcal{T}$ is modular and Lemma 1, $\mathcal{T}'$ is modular. Then $\mathscr{M}' = \langle val(\mathcal{S}'), R\rangle$ is a model of $\mathcal{T}'$, by Lemma 2. From this and Lemma 3 the result follows. ∎