Kevin Leung

Direct and Indirect Consequences of Chosen and Alternative Actions

In everyday life, much of what happens is the direct consequence of the actions we take or the events we observe. I know that it is raining because I see the rain, and I know that it's lunchtime because my roommate is calling me to lunch. Simultaneously, many things are often incidentally true or only known indirectly. For example, when my roommate calls me to lunch, I also know that it is not dinnertime. This fact is apparent from his action, yet it doesn't appear to be a direct consequence of his action. Instead, it's a consequence of the fact that he did not call me to dinner.

In Propositional Dynamic Logic (PDL), the model has a valuation that makes no distinction between direct and incidental consequences of particular actions. In this paper, I will discuss the possible extensions to PDL to capture this distinction and present a concrete application of this extension. I will also briefly discuss how this distinction might be significant to Dynamic Epistemic Logic (DEL).

The extension to PDL

In PDL, the model already "knows" everything about the consequences of actions because all worlds either validate or don't validate every formula. This new operator is intended to express more in the language of PDL itself. Specifically, we need an operator to distinguish between the consequences of choosing an action versus the consequences of having not chosen an alternative action.

Fortunately, PDL already has the operator choice, ∪, to express alternative actions. $\pi = a \cup b \cup c$ is nondeterministically choosing between a, b, and c. If action a is taken, then the alternative actions is the set of all other possible actions, being b and c in this case. Currently, [π

$\cup \pi']\varphi \leftrightarrow [\pi]\varphi \vee [\pi']\varphi$, and that will continue to be true. We need a stronger sense of choice, however, to be able to reason about the consequences of different actions. Specifically, it isn't enough to know what the alternatives are; we must also know what choice was made and what the direct consequences of the chosen action were. I propose two ways to define this action choice.

First, the choice might simply be a check afterwards to see what the last action taken was. At the risk of overloading operators, we could define actions as

$$\pi := a \mid \pi ; \pi' \mid \pi \cup \pi' \mid \varphi? \mid \pi? \mid \pi^*$$

The new action $\pi?$ is only successfully executed if π was the last action taken. For example, $(a \cup b);a?$ means that there was a choice between a and b, and a was the action taken. This operator has a very clear meaning and syntactically fits PDL. A problem , however, is that $\pi?$ needs to inspect a previous action to determine whether it executes properly and know the alternatives from the previous action. That interaction between actions seems to break some clean boundaries in the structure of the language.

Another way we might define the final choice would be bind it to $\cup$ specifically. Actions could then be defined as

$$\pi := a \mid \pi ; \pi' \mid \pi \cup \pi' \mid \varphi? \mid \pi \cup \pi', \pi'' \mid \pi^*$$

The new action $\pi \cup \pi', \pi''$ means that there was the choice between π and π', and π'' was chosen. Hopefully, π'' is either π or π', though there doesn't appear to be any immediate harm in having meaningless actions. The corresponding action to the example in the other syntax would be $a \cup b, a$. A problem with this system occurs with more than 2 choices. It's possible, in this language, to generate actions such as $((a \cup b, a) \cup c, c)$, which also shows how much trouble nesting actions with parentheses can be. Also in that example, I'm not certain how one could choose a. There's the small benefit of having a total preorder of preferences, but it does get very messy.

For the rest of this paper, I will use the $\pi \cup \pi', \pi''$ action to define the final choice between alternatives. I will carefully choose examples to avoid other complications introduced by this new action, and I hope that other syntactic issues aren't too problematic.

The model is the same as the original PDL model, $M = (S, \{R_a\}_{a \in \text{atom}}, \underline{v})$. It's worth noting again that this extension doesn't truly allow us to "know" anything more than we did before. Because the relation and valuation define all formulas in all worlds, what we know from an alternative action is already known in the resulting world itself.

The extension that we get is in our interpretation of the operator. The axiomatization for this would be

$$[\pi_1 \cup \pi_2 \cup \ldots \pi_k , \pi_1]\varphi \leftrightarrow [\pi_1]\varphi \wedge [\pi_2]\neg\varphi \wedge [\pi_3]\neg\varphi \wedge \ldots [\pi_k]\neg\varphi$$

Perhaps a simple way to read this is that some statement is a direct consequence of an action if it is the only action among alternatives that results in that statement being true. Although I don't have syntax for it, all other consequences from an action are considered indirect consequences. Those statements are the ones that are valid after this action and valid after some other action as well. If multiple actions lead to some consequence, it's less clear why exactly that comes about.

An application of this operator

Imagine that Batman has been caught in a dastardly plot by the Joker. After pursuing the Joker to his secret hideout, Batman finds Alfred, Batgirl, and Robin all suspended over vats of boiling acid. Unfortunately, Batman only has enough time to save one of them. Thus, we have 3 actions and 3 propositions: let *a* be the action of rescuing Alfred, *b* be the action of rescuing Batgirl, and *c* be the action of rescuing Robin. Let *p* be whether Alfred is alive, *q* be whether Batgirl is alive, and *r* be whether Robin is alive. The following, then, should hold if we are currently in world w:

$$M, w \models [a]\, p \wedge \neg q \wedge \neg r$$

$$M, w \models [b]\, \neg p \wedge q \wedge \neg r$$

$$M, w \models [c]\, \neg p \wedge \neg q \wedge r$$

If we only have basic PDL, perhaps the best we currently get is

$$M, w \models [a \cup b \cup c]\, (p \wedge \neg q \wedge \neg r) \vee (\neg p \wedge q \wedge \neg r) \vee (\neg p \wedge \neg q \wedge r)$$

So if Batman saves one of them, exactly one of them will be saved. To simply take one action alone without considering the other choices, though, is also a disservice to Batman as it suggests that 2 people died because of his action. At the risk of sounding political, there's a difference between not saving someone and causing them to die. AlthoCompare what we get by introducing the consequences of making a specific choice:

$$M, w \models [a \cup b \cup c, a]\, p$$

$$M, w \models [a \cup b \cup c, b]\, q$$

$$M, w \models [a \cup b \cup c, c]\, r$$

Here, we clearly distinguish between what happened because of Batman's action and what didn't happen because of other actions not taken. In this sense, $\neg q$ and $\neg r$ are peripheral to the main action. Robin is not alive because Batman saved Alfred; Robin is not alive because Batman didn't save Robin.

Thoughts on its impact on DEL

Direct and indirect consequences also matter in DEL where something is known after some action. We might interpret direct consequences as knowledge from the action itself whereas indirect consequences are knowledge known from alternatives. On the surface, the change to DEL is similar to what we have previously seen. For example, the following still holds:

$$[\pi_1 \cup \pi_2 \cup \ldots \pi_k , \pi_1]K\varphi \leftrightarrow [\pi_1]K\varphi \wedge [\pi_2]\neg K\varphi \wedge [\pi_3]\neg K\varphi \wedge \ldots [\pi_k]\neg K\varphi$$

In other words, we know something directly from an action if we wouldn't know it having taken any other action. If we represent the different worlds as event models in themselves, like in Public Announcement Logic (PAL), we can even begin to speak more concretely about what effects certain actions have on how the model changes. Take the model changes below:

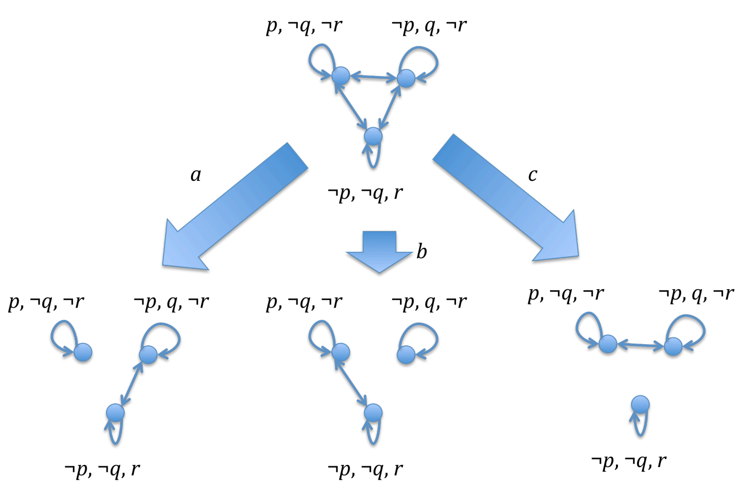

If we imagine that the upper left world is the actual world, we see that $M, w \models [a](Kp \wedge K\neg q \wedge K\neg r)$ and $M, w \models [a \cup b \cup c, a]Kp$. If we use the interpretation of a, b, c, p, q, and r from the example above, it's clear that we appear to have "lucked" into $K\neg q$ and $K\neg r$. Both of these statements are valid only because of state of the model, not necessarily because of the action itself. This difference is similar to what occurs in PAL: we may only announce φ, but ψ may also be valid in the resulting model. The addition that this operator brings, though, is that we also see other direct consequences. For example, if $\varphi \leftrightarrow \psi$, ψ will also be a direct consequence of the action, whereas an announcement only specifies the announced formula.

Discussion

Although these ideas don't seem to add new truths to PDL, they do perhaps make a finer distinction in the source of consequences and how that might come about. On a larger scale, these extensions generally address a hole in PDL: validity only through action. Although this constraint is the nature of PDL, we also accumulate knowledge through inaction or an expected action not occurring. This new operator doesn't quite address that question, but the idea of indirect consequences expands the types of knowledge that PDL can capture.

Beyond the big syntactic concerns discussed above, we also haven't explored what happens when the action chosen is more complicated than an atomic action. The new operator is defined as one-against-many: it's a direct consequence if it's not for any alternative. Situations exist where a statement is a direct consequence of a subset of potential actions, and it should be a direct consequence of those. For example, we might consider a formula such as $[\pi_1 \cup \pi_2 \cup \ldots \pi_k, \pi_1 \cup \pi_4]\varphi$, where φ is a direct consequence of π_1 and π_4, In addition to capturing direct and indirect consequences and sources of knowledge, it also includes the idea of intention in actions. As we saw in the Batman example above, my semantics for this operator only has a weak sense of intention in that only propositions that result only from this action are seen as direct consequences. The strongest interpretation is that all formulas in the resulting world are consequences of the action without consideration for alternatives, which is roughly what PDL currently has. Although intention is harder than these extreme interpretations would treat it, this new operator hopefully gives question and others an interesting direction for interpretation in PDL and DEL.