Progress Towards a Formal Theory of Practical Reasoning: Problems and Prospects

Richmond H. Thomason
AI Laboratory
130 ATL Building
University of Michigan
Ann Arbor, MI 48109–2210

`rich@thomason.org`

These pages are copies of transparencies that were used in a talk given at the IJCAI-99 Workshop on Practical Reasoning and Rationality, Stockholm, July 31, 1999. They were prepared for private distribution by email and over the Internet on August 19, 1999.

A paper on the formalism described in the second part of the talk is in preparation, and should be accessible from the author's home page, www.eecs.umich.edu/~rthomaso, by September 15, 1999.

Comments are welcome. Please email them to `rich@thomason.org`.

Abstract

From its beginnings in Aristotle, logic was intended to account not only for reasoning that is *theoretical* (or conclusion-oriented), but for reasoning that is *practical* (or action-oriented). However, despite an interest in the topic that continues to the present, the practical side of reasoning has remained broadly speculative. At least in some domains (mathematics, in particular), there are well developed proof-theoretic and semantic theories that yield quite detailed models of correct reasoning, and these models are useful for both theoretical and practical purposes. In contrast, the logical work on practical reasoning has remained broadly speculative and disengaged from applications. Logical formalisms have not been forthcoming that would be useful either in designing an agent that needs to act intelligently, or in helping an intelligent agent to evaluate its reasoning about action.

The decision-theoretic paradigm that has dominated economic thinking, on the other hand, certainly has produced applicable models of correct decision making. And, though decision theory and logic are certainly different subjects, it is easy to find areas of overlap in the concepts and techniques, as well as people who have made fundamental contributions to both fields.

Despite these similarities, I think it would be wrong to think of decision theory as the realization, within a different academic discipline, of a logical theory of practical reasoning. The reason is that *correct inference* is central to the logical approach to a subject matter, and correct inference is largely neglected in the decision theoretic paradigm.

The absence of a logical theory of practical reasoning is largely due to the unavailability of appropriate inference procedures. To handle even the simplest cases of practical reasoning, it is essential to have a reasoning mechanism that allows for practical conclusions that are nonmonotonic in the agent's beliefs. If an agent believes that he is out of milk, he may well conclude to walk to the store. If he then adds the belief that the store is closed, he will then have to withdraw his conclusion. And, until recently, probability functions have provided the only way to formalize inference procedures with these characteristics.

An approach to practical reasoning based on probability relies on numerical calculation rather than qualitative inference, so it needs quantities, not only for probabilities, but for utilities. Leonard Savage called the problem of constructing a quantitative model the *small worlds problem.* A good solution to the small worlds problem is great when you can get one. But you can't always do that. Trying to deal with decision problems in the absence of a qualitative model raises a number of difficult questions.

1. How to represent the reasoning process (rather than just the outcome).
2. How to make use of large amounts of knowledge, in open-ended decision situations. (In practice, the decision theoretic models are limited to outcomes that depend on no more than a dozen or so variables.)
3. How to make use of reasonable assumptions that are known in some sense, but cannot readily be assigned a probability in a many contexts.
4. How to construct a decision-theoretic microworld.
5. How to learn an agent's preferences from readily available information.
6. How to deal with conflicting goals.
7. How to model cases in which the agent is to some extent distributed, without complete agreement or communication among sub-agents.
8. What to do about problems of real-time, resource limited reasoning.

To deal with these problems, we need an alternative theory of a decision-making agent with the following characteristics.

Regarding belief the theory should:

1. Relax the quantitative commitments of decision theory.
2. Provide for belief kinematics, in allowing an update function to be defined.
3. Be engineerable. In particular, the information needed to support and update beliefs should be acquirable in some practicable way.

Regarding desire, the theory should:

1. Retain the idea that desires are immediate, with a source that is external to practical reasoning (below, I will call these immediate desires *wishes*), and that there are reasoned desires that depend on wishes and beliefs (below, I will call these *wants*). It is assumed that wants are like intentions, which are more or less connected to actions.
2. Treat practical reasoning as a process that creates considered desires by transforming wishes into wants.
3. Allow for the creation, cancellation, and reprioritization of wants in light of changing beliefs.
4. Treat the outcome of practical reasoning as nonunique. Agents with the same beliefs and wishes could reach different conclusions, even while conforming to the full principle of rationality.

Developing such a theory makes for a large-scale challenge. However, new ideas from many disciplines (and especially from Artificial Intelligence) provide a real opportunity for meeting this challenge.

In my talk, I will try to provide a sketch of these opportunities.

*********************** Slide 1 ***********************

Practical Syllogism

Aristotle, 4th Century B.C.

I desire something sweet.

This is sweet.

An act of tasting occurs.

Belief + Desire → Performed Action

*********************** Slide 2 ***********************

Decision Theory

Savage, Jeffrey, 20th Century A.D.

A Probability Function.

A Utility Function.

An Assignment of Expected Utilities to Actions.

Belief + Desire →

Recommended Action

************************ Slide 3 ************************

Aristotle's Distinction Between Theoretical and Practical Reason

It can be stated in terms of agent attitudes:

Theoretical Reason:

Belief → Belief

Practical Reason:

Belief + Desire → Intention

The relation between Intention and Action would be explained in the agent architecture.

************************ Slide 4 ************************

Axiomatization Versus Formalization

Axiomatization:

- **Choose a reasoning domain.**
- **Develop a declarative logic with a consequence relation.**
- **Present axioms.**
- **Show that conclusions in the domain can be obtained using the logic from axioms.**

Formalization:

- **Choose a reasoning domain.**
- **Devise Formal Representations.**

- Provide reasoning algorithms. (Hopefully, there is some apparatus for understanding the algorithms.)
- Test the adequacy of the algorithms by comparing with specimens of correct reasoning.

Algorithms can include:

- Probability Estimates
- Utility Calculations
- Scoring algorithms of various sorts, as in CBR

*********************** Slide 5 ***********************

Logicist Projects

- Theoretical Reasoning Domain
- Use Axiomatization Methodology

Reasoning Formalization Projects

- Practical Reasoning Domain
- Acquire reasoning specimens, take protocols
- Formalize
- Implement
- Compare results

Examples:

Expert Systems

Case-Based Reasoning

Cognitive Modeling

Most of NLP not included

************************** Slide 6 **************************

There is a lot of reasoning in practical reasoning

- **Manuals, instructions, recipes**
- **Legal argumentation**
- **Accusations, excuses**
- **Planning, configuring**

************************** Slide 7 **************************

Back to Aristotle's Practical Syllogism

Not really a contender as a theory of practical reasoning.

Problems with Expressive Power, similar to those of the theoretical syllogism.

Desires can be inferred, as well as beliefs.

************************ Slide 8 ************************

Decision Theory

1. **Belief:** $P(A)$, $0 \leq P(A) \leq 1$
2. **Belief kinematics:**

   $P_A(B) = P(B/A)$, **where** $P(A) \neq 0$.
   (Note: this can be generalized, using Popper functions or nonstandard probabilities.)

3. **Desire:**

   (a) **Absolute Desires:**

   $\$(w)$ defined over outcomes w.

   (b) **Belief-Relative Desires:**

   $Ut(A)$, **where** $Ut(A) = \sum_{w \in A} p(w)\$(w)$

************************ Slide 9 ************************

Problems with Decision Theory

1. **How to represent the reasoning process (rather than just the outcome).**
2. **How to make use of large amounts of knowledge, in open-ended decision situations. (In practice, the decision theoretic models are limited to outcomes that depend on no more than a dozen or so variables.)**
3. **How to make use of reasonable assumptions that are known in some sense, but cannot readily be assigned a probability in a many contexts.**
4. **How to construct a decision-theoretic microworld.**
5. **How to learn an agent's preferences from readily available information.**
6. **How to deal with conflicting goals.**
7. **How to model cases in which the agent is to some extent distributed, without complete agreement or communication among sub-agents.**
8. **What to do about problems of real-time, resource limited reasoning.**

*********************** Slide 10 ***********************

A Problem with All Formalisms?

Poor Mapping Between Reasoning Specimens and Formalized Reasoning

I need to teach summer school in Chicago.

I'd like to have a car, but it's too far to drive, and it would be too expensive to rent a car.

So I'll fly to Chicago.

I'll need some clothing and a big box of books, and a laptop computer.

It would be a hassle to get all that stuff to the airport and take it on the plane.

So I'll ship the books.

So I might as well ship all the stuff I'll need.

But the laptop is expensive—I'll carry that.

*********************** Slide 11 ***********************

A Corollary of the Difference between Theoretical and Practical Reasoning

There will be no clear, ready-made methodology.

There will be difficulties in evaluating theories.

There will be difficulties in collecting and organizing large, coherent collections of data, with long-term, multiple uses.

But in AI, cognitive modeling, and economics, methodologies have emerged that are relevant and helpful.

*********************** Slide 12 ***********************

So Much for Background, and the Big Picture

The remainder of this paper will:

- Concentrate on the parts that can be tackled with a methodology that is broadly logical.
- Illustrating opportunities for progress that have emerged from work in logical AI.

*********************** Slide 13 ***********************

Action

There is no robust, independent theory of action in decision theory. Actions are identical in type with propositions.

There is an extensive literature on action in recent philosophy, going back to the 1960s. But philosophers have concentrated on how actions fit into the causal order. There is very little in the philosophical literature that illuminates reasoning about action.

The work that has been done on planning formalisms in AI represents a big step forward. By focusing on the feasibility and effects of actions, it manages to formalize an important part of means-end reasoning.

*********************** Slide 14 ***********************

A Formalization Discipline for Actions

Note: This is a logicist program.

- Action-centered actions, in a temporal logic. (Various temporal formalisms are used for this purpose.)
- Axioms give preconditions, effects of actions.
- Basic reasoning process: prediction, conditional on action sequences.
- Qualitative, deductive approach to prediction.

Organization of action and change formalization problems

Frame Problem

Qualification Problem

Ramification Problem

Increasing emphasis on causality.

Relations to BDI agent architectures.

********************** Slide 15 **********************

Limitations of Planning Formalisms

1) **The logicist planning community doesn't recognize the importance of plan evaluation.**
2) **There is no good way to deal with plan monitoring, plan modification.**
3) **There is no way to deal with uncertainty and risk, at least within the limits of this formalism.**
4) **Goals are simply given, and on some accounts are adhered to until achieved or shown infeasible. There is no reasoning about desires.**

********************** Slide 16 **********************

"Background to Qualitative Decision Theory"

Jon Doyle and Rich Thomason
AI Magazine, Summer, 1999

Problem of sympathetic planning.

Generic Preferences, and their adaptation to specific circumstances.

*********************** Slide 17 ***********************

Suggestions about How to Generalize Some Ideas from Logical AI to Obtain a More Robust Formalism for Reasoning with Desires

Nonmonotonic Logics

Commonsense defaults can conflict.

Resolution of conflicts leads to multiple conclusion sets.

Priorities may be needed.

In terms of agent attitudes, a distinction may be needed between immediate, inferred beliefs.

These ideas are also very appropriate for formalizing desires.

*********************** Slide 18 ***********************

***Prima Facie* and All-Things-Considered Beliefs**

There are *prima facie* and all-things-considered beliefs.

Example 1. Beliefs about the porch light.

(i) **I have a reason to believe the porch light is off, because I asked my daughter to turn it off.**

(ii) **I have a reason to believe the porch light is on, because the last time I saw it, it was on.**

(iii) **All things considered, I believe the porch light is off, because my daughter is pretty reliable.**

*********************** Slide 19 ***********************

Some Considerations

***Prima facie* beliefs (i.e., reasons for believing something) can conflict with one another.**

In deciding what we believe, all-things-considered, we need to reconcile such conflicts, if they exist.

Like an intention, the belief that the light is off acts as a constraint on future deliberation.

Nonmonotonic logics which, like Default Logic, involve construction and selection of extensions, provide a model of this sort of reasoning process. *Prima facie* beliefs are like default rules and consequences that can be derived from these defaults and information that is already established. All-things-considered beliefs are like the component propositions of a single extension that has been selected as a consequence of the default rules.

************************ Slide 20 ************************

Wishes/Wants

Part of commonsense practical reasoning consists in the practicalization of desires. Immediate desires needn't be feasible, and typically will conflict with other immediate desires. We do not expect all of these wishes to survive as practical goals. The ones that do survive I will call *wants.*

This distinction seems to correspond to one important difference beteen the way 'wish' and 'would like' on the one hand and 'want' on the other are typically used. In particular:

Wishes can conflict with beliefs.

> *Example 2.*
> I'd like to take a long vacation.
> I'd need to get time off from work to take a long vacation.
> *But:* I can't get time off from work.

Wishes can conflict with each other, in lignt of background beliefs.

> *Example 3.*
> I'd like to take a long vacation.
> *But:* I'd like to save more money this year.
> *And:* I can't save more money this year and take a long vacation.

Wishes can conflict with intentions, or more generally with adopted plans. This point is made by Michael Bratman, David Israel, and Martha Pollack. See [2, 1]. For present purposes, it is not important to distinguish between wants and intentions.

*********************** **Slide 21** ***********************

An Analogy We Can Exploit

Wishes are like *prima facie* beliefs.

So, use defaults for both. But use a notation that lets us keep track of which is which.

$$A \overset{\mathbf{B}}{\hookrightarrow} C$$

versus

$$A \overset{\mathbf{D}}{\hookrightarrow} C.$$

(Note: we are limiting ourselves to normal discourse.)

Wants are like all-things-considered beliefs.

So treat both as conclusions in a selected extension generated by the defaults. Don't distinguish the two types of conclusions notationally. But we can account for the difference in terms of the reasons that explain why a conclusion belongs to the extension.

*********************** **Slide 22** ***********************

A Reasoning Example

***Example 4. Part I:* Commonsense reasoning.**
(Imagine a restaurant scenario.)

1. **I'd like to have some coffee.**
2. **For me to have coffee, coffee will have to be available.**
3. **I'd like to have decaf if I have coffee.**
4. **Defaf must be available if coffee is available.**
5. **Coffee is available.**
6. **For me to have decaf coffee, I'll need to order decaf coffee.**
7. ***So:* I'll order decaf coffee.**

*********************** Slide 23 ***********************

A Reasoning Example, Continued

Example 4. Part II: **Formalization.**

$\top \overset{D}{\hookrightarrow}$ Coffee
Coffee $\overset{B}{\hookrightarrow}$ Available
Coffee $\overset{D}{\hookrightarrow}$ Decaf
Available $\overset{B}{\hookrightarrow}$ Decaf-Available
Available
Decaf $\overset{B}{\hookrightarrow}$ Order-Decaf

There is one extension, which is generated by the following choices:

$$\{\text{Coffee}, \text{Available}, \text{Decaf}, \text{Decaf-Available}, \text{Order-Decaf}, \}$$

Note: The use of the premise Decaf $\overset{B}{\hookrightarrow}$ Order-Decaf is a makeshift. The selection of an action to achieve an end should be carried out by means of a planning process. I intend to explain in a later paper how to integrate the formalism with planning.

*********************** Slide 24 ***********************

Another Reasoning Example

Example 5. Part I: **Commonsense reasoning.**
(Imagine a hiking scenario.)

I think it's going to rain.

If it rains, I'll get wet.

(Even) if it rains, I wouldn't like getting wet.

*********************** Slide 25 ***********************

Another Reasoning Example, Continued

Example 5. **Part II: *Formalization.***

$\top \overset{\mathbf{B}}{\hookrightarrow}$ **Rain**
Rain $\overset{\mathbf{B}}{\hookrightarrow}$ **Wet**
Rain $\overset{\mathbf{D}}{\hookrightarrow} \neg$**Wet**

Intuitively, there should be only one extension, which is generated by the following choices:

$$\{\mathbf{Rain}, \mathbf{Wet}\}$$

Explanation: **If I genuinely believe that it will rain, and that I will get wet if it rains, I should believe that I will get wet, regardless of my preferences or likings. To do otherwise would be to indulge in *wishful thinking.***

Note that wishful thinking is incorrect, regardless of whether the beliefs are defeasible. To take another example, suppose that I am not at home at t, that my umbrella is home at t, and that I would like to have my umbrella at $t+1$. If there were no prohibition of wishful thinking, nothing would prohibit a plan in which at t I simply wait, achieving my goal by wish fulfillment.

To put it another way, if practical reasoning is to be practical, wishes can't be fulfilled simply because they are wishes, but have to be achieved by feasible actions. It is belief that determines feasibility, not desire. So in a formalism that allows desire-based defaults, belief-based defaults have to take precedence.

In the context of the formalism I am proposing, prioritization of belief over desire can be achieved by using a standard account of prioritized defaults. I wlll not discuss the technical details here. See, for instance, [3].

With belief defaults prioritized over desire defaults, we obtain a single extension in Example 5, the one generated by the following choices.

$$\{\mathbf{Rain}, \mathbf{Wet}\}$$

*********************** Slide 26 ***********************

Still Another Reasoning Example

Example 6. Part I: **Commonsense reasoning.**
(Imagine a hiking scenario. This example is more complex than the others. There will be multiple extensions.

I think it will rain.

If it rains, I'll get wet.

(Even) if it rains, I wouldn't like to get wet.

If I get wet, I'd like to change into dry clothes.

If I change into dry clothes, I'll have to walk home.

If I walk home, I'll have to walk an extra two hours.

I wouldn't like to walk an extra two hours.

*********************** Slide 27 ***********************

Still Another Reasoning Example, Continued

Example 7. Part II: **Formalization.**

$\top \stackrel{\mathbf{D}}{\hookrightarrow} \mathbf{Rain}$
$\mathbf{Rain} \stackrel{\mathbf{B}}{\hookrightarrow} \mathbf{Wet}$
$\mathbf{Rain} \stackrel{\mathbf{D}}{\hookrightarrow} \neg\mathbf{Wet}$
$\mathbf{Wet} \stackrel{\mathbf{B}}{\hookrightarrow} \mathbf{Change}$
$\mathbf{Change} \stackrel{\mathbf{B}}{\hookrightarrow} \mathbf{Home}$
$\mathbf{Home} \stackrel{\mathbf{B}}{\hookrightarrow} \mathbf{Walk\text{-}Two\text{-}Hours}$
$\top \stackrel{\mathbf{B}}{\hookrightarrow} \neg\mathbf{Walk\text{-}Two\text{-}Hours}$

There are two extensions.

The first is generated by the following choices:

$$\{\mathbf{Rain}, \mathbf{Wet}, \mathbf{Change}, \mathbf{Home}, \mathbf{Walk\text{-}Two\text{-}Hours}\}$$

The second is generated by the following choices:

$$\{\text{Rain}, \text{Wet}, \text{Change}, \neg\text{Walk-Two-Hours}\}$$

These two extensions represent the two choices that the scenario of Example 7 makes available; on the one hand getting dry, but walking two hours, and on the other staying wet, but avoiding the extra walk.

In a purely epistemic version of default logic, multiple extensions represent equally reasonable alternatives, and there is no way to choose between them. With desire-based defaults added to the mix, multiple extensions are still in a sense equally reasonable, but these are the sorts of choices that numerical utilities are designed to resolve, and we may be able to choose between extensions using, for instance, a multi-attribute utility analysis. An agent who strongly dislikes being wet, but who enjoys walking, for instance, will prefer the second extension.

*********************** Slide 28 ***********************

Some Concluding Thoughts

In thinking about practical reasoning, we should keep the big picture in sight.

We can't really expect the broader problems of practical reasoning to succumb to a narrow logical effort.

But trends in AI, economics, and cognitive science can contribute.

We need more collaborations that bring these trends together.

At the same time, we may be able to move ideas based on the logical paradigm in a direction that could help the larger effort.

************************ Slide 29 ************************

References

[1] Michael Bratman. *Intentions, Plans and Practical Reason.* Harvard University Press, 1987.

[2] Michael Bratman, David Israel, and Martha Pollack. Plans and resource bounded practical reasoning. *Computational Intelligence*, 4:349–355, 1988.

[3] Gerhardt Brewka. Reasoning about priorities in default logic. In Barbara Hayes-Roth and Richard Korf, editors, *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 940–945, Menlo Park, California, 1994. American Association for Artificial Intelligence, AAAI Press.