

The Formalization of Practical Reasoning: An Opinionated Survey*

Richmond H. Thomason
University of Michigan

The Challenge

I begin by considering examples of practical reasoning. In the remainder of the paper, I try to say something about what a logical approach that begins to do justice to the subject might be like.

Example 1. Ordering a meal at a restaurant. The deliberating agent sits down in a restaurant and is offered a menu. Here, the problem is deciding what to eat and drink. Even if the only relevant factors are price and preferences about food, the number of possible combinations is very large. The reasoning may involve weighing preferences about food and drink against preferences about cost, but might well produce a decision without a general rule for reconciling all the preferences.

Example 2. Deciding what move to make in a chess game. A classic AI example.

Example 3. Savage’s omelet. Five good eggs have been broken into a bowl. A sixth egg is unbroken. The example involves preferences about the desired outcomes, as well as risk, in the form of a positive probability that the egg is spoiled. The task is to infer preferences over actions. The outcomes involve only a few variables, the preferences over them are evident, and the probabilities can be easily estimated. The reasoning reduces to the calculation of an expected utility. (Savage 1972)[pp. 13–15].

Example 4. Designing a house. This example is less obviously practical; it is possible for an architect to design a house without thinking much about the actions that will go into building it, leaving this to the contractor.

Example 5. Deciding how to get to the airport. Another classic AI example.

Example 6. Cracking an egg into a bowl. A classic common sense reasoning example.

Example 7. Playing table tennis. Here, practical reasoning has to be engaged in complex, real-time activities involving the perceptual and motor systems. There is no

time to spare for reflection; the reasoning needs to be thoroughly connected to the ongoing process of play.

Example 8. Playing soccer. Soccer is like table tennis, but with the added dimension of teamwork and the need to recognize and execute play. This task was selected as a benchmark problem in robotics, and has been extensively studied. A classic robotics example.

Example 9. Typing a message. Typing an email message, composing it as you go along, starts perhaps with a general idea of what to say. The reasoning that produced a rough idea of the content may have taken place reflectively, but once composition has begun, several reasoning processes are engaged simultaneously, and have to be coordinated. The general idea of what to say has to be packaged in linguistic form, which must be rendered by motor actions at the keyboard. For a skilled typist composing a straightforward message, these complex, practical tasks are combined and executed very quickly, perhaps at the rate of 70 words per minute. For this to happen, the interface between high-level linguistic reasoning and motor skills has to be very robust.

Example 10. Factory scheduling. Another classic AI example. Part of its interest lies in the difference in scale between this problem and Savage’s omelet problem. It is not clear that there is any way to construct a single, coherent utility function for the task, by reconciling the four desiderata mentioned above. Any reconciliation will leave some managers unhappy: salesmen will favor some goals and production managers will favor others. Nor is it easy to produce a global probability function for a system with so many interacting variables.

Example 11. Ordering dessert. Our agent has a craving for dessert. This colors the alternative emotionally. But suppose that there is a contrary emotion. The agent is unhappy with being overweight and has determined to eat less, and may have told others at the table about the decision to undertake a diet. This creates a conflict, coloring the choice of dessert with negative associations, perhaps even shame. The chief difference between this conflict and those in Examples 2 and 4 is that this decision is emotionally “warm;” the outcome may be influenced by a craving and the presence of the desired object. (Perhaps this is why some restaurants invest in dessert trays.)

Copyright © 2011, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

*This is an abbreviated version of a longer paper.

¹Of course, a good design has to take into account how to build a house, in order to make sure that the design is feasible.

Example 12. An unanticipated elevator. A man decides to visit his stockbroker in person, something he has never done. He takes a bus to a stop near the stockbroker's downtown address, gets off the bus, locates the building and enters it. He finds a bank of elevators, and sees that the stockbroker is on the 22nd floor. This man has a strong dislike for elevators, and is not feeling particularly energetic that day. He reconsiders his plan.

Example 13. A woman is working in her garden. She becomes hot and tired, and decides to take a break. Or she hears the telephone ringing in her house, and decides to answer it. Or she sees smoke coming out of the window of her house, and runs for help.

Example 14. The wrath of Achilles. In Book I of *The Iliad*, the hero Achilles is outraged and dishonored by his warlord Agamemnon, who insults him and declares that he will take back the captive woman that Achilles had received as his war prize. Achilles is almost overcome with rage, but at the last moment draws back from killing the king. Even though it is strongly informed by emotion ("hot"), there is reasoning between the initial emotional storm and the action.

Example 15. Conversation. Conversation provides many good examples of deliberative reasoning. Where there is conscious deliberation, it is likely to be devoted to content selection. But the reasoning that goes into deciding how to express a given content can be quite complex.

Certainly, any adequate theory of practical reasoning must at least be compatible with this broad range of cases. Better, it should be capable of saying something about the reasoning involved in all of them. Even better, there should be a single architecture for practical reasoning, capable of dealing with the entire range of reasoning phenomena. No doubt, there are special-purpose cognitive modules (e.g., for managing perception, motor behavior, and some aspects of language). But it would be perverse to formulate a theory of a special type of practical reasoning, such as preference generation, probability estimation, or means-end reasoning, and to postulate a "cognitive module" that performs just this reasoning. This methodology would be likely to produce an *ad hoc* and piecemeal account of practical reasoning.

These examples suggest a set of features that can be used to classify specimens of deliberative reasoning.

1. Are only a few variables (e.g., desiderata, causal factors, initial conditions) involved in the decision?
2. Do conflicting preferences need to be resolved in making the decision?
3. Is the time available for deliberation small compared to the time needed for adequate reflection?
4. Is the deliberation immediate? That is, will the intentions that result from the deliberation be carried out immediately, or postponed for future execution?
5. Is the deliberation carried out in "real time" as part of an ongoing activity involving sensory and motor activities?
6. Does the reasoning have to interface closely with sensory and motor systems?
7. Is the activity part of a group or team?
8. Does the context provide a definite, relatively small set of actions, or is the set of actions open-ended?

9. Is there certainty about the objective factors that bear on the decision?
10. Is the associated risk small or great?
11. Is the goal of deliberation a single action, or a sequence of actions?
12. Is continuous time involved?
13. Is the deliberation colored with emotions?
14. Is the action habitual, or automatic and unreflective?
15. Is there conscious deliberation?
16. Are there existing plans in play to which the agent is committed or that already are in execution?

There is nothing wrong with concentrating on part of the space delineated by these features to see what can be learned from it. Chess and decision problems that, like Savage's omelet, involve a solution to the "small worlds problem," provide good examples of cases where this methodology has paid off. But to concentrate on these cases without paying any attention to the broad spectrum of examples runs the risk of producing a theory that will not contribute usefully to something more general.

Disciplines and approaches

Many different disciplines have something to say about practical reasoning. The main theoretical approaches belong to one of the five following areas: (1) Philosophy, (2) Logic, (3) Psychology, (4) Decision Theory and Game Theory, (5) Artificial Intelligence.

Philosophy

The topic of practical reasoning goes back to Aristotle. In the Twentieth Century there was a brief revival of philosophical interest in the topic. This coincided more or less with early work on deontic and imperative logic, and was carried out by a group of logically minded philosophers and a smaller group of philosophically minded logicians. In the 1960s, a group of British philosophers became interested in the topic. This period saw 10 or more articles relevant appearing in journals like *Analysis*. Of these, (Kenny 1966) seems to have the most interesting things to say about the problem of formalizing practical reasoning.²

Kenny begins with Aristotle's practical syllogism, taking several specimens of means-end reasoning from the Aristotelian corpus, and beginning with the following example.

Example 16. A doctor prescribing.

This man is to be healed.
 If his humors are balanced, he will be healed.
 If he is heated, his humors will be balanced.
 If he is rubbed, he will be heated.
 So I'll rub him.

The premisses of the reasoning, according to Kenny, are either (i) desires or duties, or (ii) relevant facts. And he characterizes the conclusion as an action. Kenny points out that this sort of reasoning doesn't fit Aristotelian syllogistic, and that a straightforward modern formalization of it would be invalid. To put it crudely, the inference from $P, Q \rightarrow P, R \rightarrow Q$, and $S \rightarrow R$ to S is invalid. Kenny has indicated an important form of practical reasoning, and pointed out a glaring problem with the propositional calculus as a formalization medium. But he had no good positive proposal.

Another trend, which tried to absorb imperative and practical inference into some sort of modal logic, was also underway in the 1960s. (Lemmon 1965) provides a logic of imperatives that prefigures the STIT approach of (Belnap, Jr., Perloff, and Xu 2001), hence a modal approach that brings in the idea of causing a state of

²For more about this period, see (Green 1997).

affairs. But reasoning in deontic logic is deductive, and if you formalize typical specimens of means-end reasoning like Example 16 in these systems, the formalizations will be invalid.

In retrospect, we can identify several assumptions that made this attempt to formalize practical reasoning unsuccessful.

1. These philosophers relied too much on deductive inference, with the propositional calculus as a paradigm, and too little on models;
2. They worked with overly simple formal languages;
3. They didn't formalize actions explicitly;
4. They missed the insight that means-end reasoning is more like abduction or heuristic search than deduction.

More recent philosophical work on practical reasoning hardly touches on the formalization problem, and tends to concentrate on a few simplistic examples.

Logic

Georg Henrik von Wright, like Kenny, begins with Aristotle's practical syllogism. But he avoids the problem of invalidity by changing the premisses. For instance, 'If he is heated, his humors will be balanced' in Example 16 becomes 'Unless he is heated, his humors will not be balanced'. This makes it easier to formulate the inference in a deontic logic, and to see how the formalization might be valid. At the same time, it is more difficult to imagine that the premisses are true. In this example, for instance, there is surely more than one way to heat the patient.

Von Wright proposes modal logic, and in particular deontic logic, as the formalization medium for practical reasoning. He also characterizes his version of deontic logic as a "logic of action," but all this seems to mean is that the atomic formulas of his language may formalize things of the form 'Agent A does action a.' He has little or nothing to say about reasoning about action.

The STIT approach to agency provides a model-theoretic account of how actions are related to consequences that is quite different from the ones that emerged from the attempts in AI to formalize planning. The connections of STIT theory to practical reasoning are tenuous, and I will not have much to say about it.

Philosophy and philosophical logic have served over the years as a source of ideas for extending the applications of logic, and developing logics that are appropriate for the extensions. I would very much like to see philosophy continue to play its foundational and creative role in developing new applications of logic, but I don't see how this can happen in the area of practical reasoning unless philosophers study and assimilate the recent contributions of computer scientists.

Psychology

As early as 1947, Herbert Simon noted divergences between decision-making in organizations and the demands of ideal rationality that are incorporated in decision theory. An important later trend that began in psychology, with the work of Amos Tversky and Daniel Kahneman, studies these differences in more detail, providing many generalizations about the way people in fact make decisions and some theoretical models.

All this raises a challenging foundational problem, one that philosophers might be able to help with, if they gave it serious attention. What level of idealization is appropriate in a theory of deliberation? What is the role of "rationality" in this sort of idealization?

Decision Theory and Game Theory

Research in behavioral economics has made microeconomists generally aware that, in their original and extreme form, the idealizations of decision theory don't account well for a broad range of nat-

urally occurring instances of practical reasoning. Attempts to mechanize decision-making led computer scientists to much the same conclusion.

A natural way to address this problem begins with decision theory in its classical form and attempts to relax the idealizations. Simon made some early suggestions along these lines; other, quite different proposals have emerged in AI and psychology. These ideas are perfectly compatible with what I will propose here. A general account of practical reasoning has to include calculations that somehow combine probability (represented somehow) and utility (represented somehow), in order to estimate risk. The more adaptable these methods of calculation are to a broad range of realistic cases, the better. But there is more to practical deliberation than calculation; calculation needs to be integrated with inferential reasoning.

In many cases of practical reasoning, conflicts need to be identified and removed or resolved. Work by economists on value trade-offs is relevant and useful here.

Computer science and artificial intelligence

For this audience, I can skip most of the details in this section. The trends in AI that I think are most important for developing a formalization of practical reasoning are: (i) means-end reasoning, (ii) reasoning about preferences, and (iii) agent architectures.

Dynamic logic and imperative inference. When an agent is given instructions and intends to carry them out unquestioningly, there is still reasoning to be done, and the reasoning is practical—although, as the instructions become more explicit, the less scope there is for interesting reasoning from the human standpoint. Even so, the logic of computer programs, where explicitness has to be carried out ruthlessly, can be instructive, because it shows how logical theory can be useful, even when the reasoning paradigm is not deductive. This is dynamic logic.

Inference, in the form of proofs or a model theoretic logical consequence relation, plays a small part in the theory of dynamic logic. Instead, *execution* is crucial. This idea is realized as the series of states that the agent (an idealized computer) goes through when, starting in a given initial state, it executes a program. Because states can be identified with assignments to variables, there are close connections to the familiar semantics of first-order logic.

Dynamic logic has led to useful applications and has made important and influential contributions to logical theory. It is instructive to compare this to the relatively sterile philosophical debate concerning "imperative inference" that took place in the 1960s and early 1970s. Three factors seem to have rendered the earlier debate about imperative inference unproductive:

1. Too great a reliance on deductive paradigms;
2. Leaving the executing agent out of the picture;
3. Confining attention to simple examples.

In dynamic logic, the crucial semantic notion is the correctness of an imperative with respect to a specification. Interesting examples of correctness are not likely without a formalized language that allowing complex imperatives, and without examples of imperatives that are more complicated than 'Close the door'. (The first example that is presented in (Harel, Kozen, and Tiuryn 2000) is a program for computing the greatest common divisor of two integers; the program uses a *while*-loop.) A model of the executing agent is essential.

The activity of interpreting and slavishly executing totally explicit instructions is a pretty trivial form of practical reasoning. But a logic of this activity is at least a start. In seeking to formalize practical reasoning, we should be mindful of these reasons for the

success of dynamic logic, seeking to preserve and develop them as we investigate more complex forms of practical reasoning.

Planning and the formalization of means-end reasoning.

Perhaps the most important contribution of AI to practical reasoning is the formalization of means-end reasoning, along with appropriate logics, and an impressive body of research into the meta-mathematical properties of these logics, and their implementation in planning systems. I do not need to discuss these developments here.

Reasoning about preferences It is hard to find AI applications that don't involve making choices. In many cases, it's important to align these choices with the designer's or a user's preferences. Implementing such preference-informed choices requires (i) a representation framework for preferences, (ii) an elicitation method that yields a rich enough body of preferences to guide the choices that need to be made, and (iii) a way of incorporating the preferences into the original algorithm.

Any attempt to extract the utilities needed for even a moderately complex, realistic decision problem will provide motives for relaxing the classical economic models of utility; but the need for workable algorithms seems to sharpen these motives. See (Goldsmith and Junker 2008) for examples and details, and (Doyle 2004), which provides a wide-ranging foundational discussion of the issues, with many references to the economics literature.

Of the relaxations of preference that have emerged in AI, *Ceteris Paribus* Preference Nets are one of the most widely used formalisms. Acyclic CP-nets support a variety of reasoning applications (including optimization), and can help to provide an approach to preference-based planning. See (Baier and McIlraith 2008) for details and further references. And in many realistic cases it is possible to extract the information needed to construct a CP-net.

There are extensions of this formalism that allow for a limited amount of reasoning about the priorities of features in determining overall preferences; see (Brafman, Domshlak, and Shimony 2006).

Like decision analysis, work in AI on preferences tends to concentrate on extracting preferences from a user or customer. Thinking about practical reasoning, however, produces a different emphasis. Examples 1, 4, 10, and 11 were designed to show that preferences are not automatically produced by the environment, by other agents, by the emotions, or by a combination of these things. We deliberate about what is better than what, and preferences can be the outcome of practical reasoning. The status of an agent trying to work out its own preferences, and of a systems designer or decision analyst trying to work out the preferences of a person or an organization, may be similar in some ways, but I don't think we can hope that they are entirely the same. Nevertheless, insights into methods for extracting preferences from others might be helpful in thinking about how we extract our own preferences.

Agent architectures. When the planning agent is equipped with means of gathering its own information, perhaps by means of sensors, and is capable of performing its own actions, the agent will need to perform a variety of cognitive functions, and to interleave cognitive performances with actions and experiences.

1. Many of the agent's original goals may be conditional, and these goals may be activated by new information received from sensors. This is not full autonomy, but it does provide for new goals that do not come from a second party.
2. Some of these new goals may be urgent; so the agent will need to be interruptible.
3. It must commit to plans—that is, it must form intentions. These intentions will constrain subsequent means-end reasoning, since conflicts between its intentions and new plans will need to be identified and eliminated.
4. It will need to schedule the plans to which it has committed.
5. It will need to monitor the execution of its plans, to identify flaws and obstacles, and repair them.

These needs led to BDI architectures. (See (Bratman, Israel, and Pollack 1988).) This work stressed the importance of intentions, and the role that they play in constraining future planning. Work in “cognitive robotics” provides a closely related, but somewhat different approach to agent architectures.

Developments in philosophical logic and formal semantics have provided logics and models for propositional attitudes. Using these techniques, it is possible to formulate a metatheory for BDI agency. Such a metatheory is not the architecture; the reasoning modules of a BDI agent and overall control of reasoning has to be described procedurally. But the metatheory can provide specifications for some of the important reasoning tasks. Wooldridge's logic of rational agents, *LORA*, develops this idea; see (Wooldridge 2000).

A final word. Practical reasoning is not quite the same as logici-
cist AI, or even the logical theory of BDI agents. But the successful use of logical techniques in this area of AI provides encouragement for a logical approach to practical reasoning, as well as a model for how to proceed.

Towards a formalization

The challenge is this: how to bring logical techniques to bear on practical reasoning, and how to do this in a way that is illuminating, explanatory, and useful? In this chapter, I will only try to provide an agenda for addressing this challenge. The agenda divides naturally into subprojects. Some of these subprojects can draw on existing work, and especially on work in AI, and we can think of them as well underway or even almost completed. Others are hardly begun.

Relaxing the demands of formalization

Let's return to the division between theoretical and practical reasoning.

Traditionally, theoretical reasoning domains are formalized using an explicit syntax, a model-theoretic consequence relation, and perhaps a proof procedure. There was no model of the reasoning agent, except perhaps, in the highly abstract form of a Turing machine, which would be guaranteed whenever the consequence relation is recursively enumerable. When it comes to practical reasoning, I believe that we have to be prepared to relax this picture.

My proposal is this: (1) we need to add a model of the reasoning agent, (2) we need to identify different phases of practical reasoning in agent deliberation, and different ways in which logic might be involved in each phase of the reasoning, and (3) consequently, we need to be prepared to have a logical treatment that is more pluralistic and less unified.

Agent architectures and division of logical labor

How should we model an agent that is faced with practical reasoning problems? I have suggested that we should aim at, or at least acknowledge the existence of, a very broad range of reasoning problems. Suppose, for instance, that we classify the types of

reasoning that we may need to consider in terms of the sort of conclusion that is reached. In view of the examples we began with, we will need to be prepared for the agent to infer:

1. Goals, which then invoke planning processes;
2. Plans, and the subgoals or means that emerge from plans;
3. Preferences emerging from reasoning about tradeoffs and risk;
4. Intentions, commitments about what to do, and (to an extent) about when to do it;
5. Immediate decisions about what plan to execute;
6. Immediate, engaged adjustments of ongoing activities and plan executions, and shifts of attention that can affect the task at hand.

The examples were chosen, in part, to illustrate these activities. These sorts of deliberation are distinct, and all are practical. Although some of them can be automatic, they all can involve deliberate reasoning.

These six activities comprise my (provisional) division of practical reasoning into subtasks, and of the deliberating agent into subsystems. Each of them provides opportunities for logical analysis and formalization.

Means-end reasoning

This is the best developed of the six areas. We can refer to the extensive AI literature on planning and means-end reasoning not only for well developed logical theories, but for ideas about how this deliberative function interacts with the products of other deliberative subsystems—for instance, with preferences, and with plan monitoring and execution.

The practicalization of desires

On the other hand, work in AI on means-end reasoning, and on BDI agents, has little or nothing to say about the emotions and the origins of desires. In general, it is assumed that these come from a user—although the goals may be conditional, so that they are only activated in the appropriate circumstances. In principle, there is no reason why goals couldn't be inferred or learned. But the relevant reasoning processes have not, as far as I know, been formalized.

In truly autonomous agents some desires—perhaps all—originate in the emotions. Although a great deal has been written about the emotions, it is hard to find work that could find a useful purpose of logic, although computational work on emotional agents might be useful.

However desires originate, although they may be emotionally colored, they may not all be emotionally “hot;” and, to be useful in reasoning, some desires must be conditional, and self-knowledge about conditional desires must be robust. My preference for white wine this evening will probably be colored by feelings of pleasure when I think about the refreshing taste of white wine. But the feeling of hypothetical pleasure is relatively mild; I am certainly not carried away by the feeling. And AI systems builders are interested in obtaining a large body of conditional preferences from users because preferences need to be brought to bear under many different circumstances, so that a user's unconditional preferences—the preferences that are activated in the actual state of affairs—will not be very useful. Fully autonomous agents need conditional preferences as well, in planning future actions and in contingency planning.

Perhaps—to develop the example of preference for white wine a bit further—the only mechanism that is needed to generate conditional desires is the ability to imagine different circumstances, together with the ability to color these circumstances as pleasant (to some degree), and unpleasant (to some degree). But it is unlikely to be this simple, because pleasantness is not monotonic with respect

to information: I find the idea of a glass of white wine quite pleasant, but the idea of a glass of white wine with a dead fly in it quite unpleasant. Also, my feelings about some imagined situations can be mixed, with elements that I find pleasant and elements that I find unpleasant. At this point, I might have to invoke a conflict resolution method that has little or nothing to do with the emotions.

This leads to a further point: there is a difference between raw or immediate desires, or *wishes*, and all-things-considered desires, or *wants*. This is because desires can conflict not only with one another, but with beliefs. And, when they conflict with beliefs, desires must be overridden: to do otherwise would be to indulge in wishful thinking.

In (Thomason 2000), I explored the possibility of using a non-monotonic logic to formalize this sort of practicalization of desires. I'm not altogether happy with that theory, but I still believe that the practicalization of desires is an important part of practical reasoning that provides opportunities for using logic to good advantage.

Intention formation

The product of successful means-end deliberation will be an intention, taking the form of commitment to a plan. But the deliberation would not get started without a goal—and I see no difference between a goal and a (perhaps very general and sketchy) intention. Often, even in human agents, these goals come from habits, or from compliantly accepted instructions from other agents.

But sometimes goals arise internally, as outcomes of deliberation. If the conclusion of the reasoning is a practicalized desire, say, to order the trout and a bottle of the 09 Chablis, commitment to the conclusion will produce an intention, which may even become a goal for means-end reasoning. (“How am I to place my order?”)

This is why practicalization can be an important component of practical reasoning, especially if the reasoner is an autonomous human being.

What to do now?

There will come moments in the life of an autonomous agent when there is scope for new activities. These opportunities need to be recognized, and an appropriate task needs to be selected for immediate execution. A busy agent with many goals and a history of planning may have an agenda of tasks ready for such occasions; but even so, it may take reasoning to select a task that is rewarding and appropriate. I do not know if any useful work has been done on this reasoning problem.

Scheduling, execution and engagement

The examples show that there can be deliberation even in the execution of physically demanding, real-time tasks. And there can be such a thing as overplanning, since the plans that an agent makes and then performs itself will need to be adjusted to circumstances.

Also, not all intentions are immediate. Those that are not immediate need to be invoked when the time and occasion are right.

There has been a great deal of useful work on these topic in AI; just one one recent example is (Fritz 2009).

Framing a practical problem

Leonard Savage's “Small worlds problem” is replicated in the more qualitative setting of means-end deliberation. A means-end reasoning problem requires (at least) a set of actions, a description of the initial conditions, and a goal. But, even in complex cases, formulations of planning problems don't include every action an agent might perform, or every fact about the current state of the world. Somehow, a goal (like “getting to the airport”) has to suggest a method of distinguishing the features of states (or “fluents”) and the actions that are relevant and appropriate.

References

- Baier, J. A., and McIlraith, S. A. 2008. Planning with preferences. *The AI Magazine* 29(4):25–36.
- Belnap, Jr., N. D.; Perloff, M.; and Xu, M. 2001. *Facing the Future: Agents and Choices in Our Indeterminist World*. Oxford: Oxford University Press.
- Brafman, R. I.; Domshlak, C.; and Shimony, S. E. 2006. On graphical modeling of preference and importance. *Journal of Artificial Intelligence Research* 25:389–424.
- Bratman, M. E.; Israel, D.; and Pollack, M. 1988. Plans and resource-bounded practical reasoning. *Computational Intelligence* 4:349–355.
- Doyle, J. 2004. Prospects for preferences. *Computational Intelligence* 20(2):111–136.
- Fritz, C. 2009. *Monitoring the Generation and Execution of Optimal Plans*. Ph.d. dissertation, University of Toronto, Toronto.
- Goldsmith, J., and Junker, U. 2008. Preference handling for artificial intelligence. *The AI Magazine* 29(4):9–12.
- Green, M. 1997. The logic of imperatives. In Craig, E., ed., *The Routledge Encyclopedia of Philosophy*. New York: Routledge. 717–21.
- Harel, D.; Kozen, D.; and Tiuryn, J. 2000. *Dynamic Logic*. Cambridge, Massachusetts: The MIT Press.
- Kenny, A. J. 1966. Practical inference. *Analysis* 26(3):65–75.
- Lemmon, E. J. 1965. Deontic logic and the logic of imperatives. *Logique et Analyse* 8:39–71.
- Savage, L. 1972. *The Foundations of Statistics*. New York: Dover, 2 edition.
- Thomason, R. H. 2000. Desires and defaults: A framework for planning with inferred goals. In Cohn, A. G.; Giunchiglia, F.; and Selman, B., eds., *KR2000: Principles of Knowledge Representation and Reasoning*, 702–713. San Francisco: Morgan Kaufmann.
- Wooldridge, M. J. 2000. *Reasoning about Rational Agents*. Cambridge, England: Cambridge University Press.