Reprinted from *Technical Papers of the ARPA Planning Initiative Workshop*, Tucson, AZ, February, 1994.

A Reasoning Economy for Planning and Replanning

Jon Doyle
Massachusetts Institute of Technology
Laboratory for Computer Science
545 Technology Square
Cambridge, Massachusetts 02139
Tel: (617) 253-3512, Fax: (617) 258-8682
`doyle@lcs.mit.edu`

Abstract

Major military operations and other large-scale activities naturally involve competitions for resources that must be managed effectively to ensure overall success. Since plan revisions may change the resource demands of the activity, the relevant competitions for resources involve computational resources (time, database control, etc.) as well as the more obvious non-computational resources (fuel, aircraft, etc.). Ideally, these conflicting demands should be resolved rationally in the sense of decision theory and economics. We determine rational allocations over the full range of resources by using an artificial market economy (implemented as RECON, the Reasoning ECONomy) to determine prices or trading ratios among resources. We represent tasks and goals as resource-endowed consumers and computational methods, informational resources, and reasoning procedures as resource-transforming producers. We then use the information provided by such problem-oriented market economies to guide search, planning, and revision procedures that automatically extend and revise the market structure itself, e.g., to include markets for solutions to new tasks and subtasks. This permits the market information to guide the choices of whether to revise plans, and which portions to revise if revisions are warranted.

1 Rational coordination of activities

Large-scale activities, such as major military operations or the operations of major healthcare and manufacturing organizations, naturally involve competitions for resources. In the first place, these activities usually involve several organizations or authorities, each placing demands on its own and other resources. In the second place, these competing organizations often generate internal competitions as well, as they distribute information and authority, both geographically (different theaters of operation or manufacturing or clinic locations) and functionally (special-purpose databases and departments), in order to satisfy legal, regulatory, privacy, reliability/redundancy, or efficiency (e.g., communication or computational) considerations.

The overall success of large-scale activities often requires managing these competing demands effectively to coordinate the component efforts, and recognizing that computational and communication resources must be allocated along with fuel, diagnostic equipment, repair personnel, and other more visible goods. Computation and communication may seem "free" once one makes these resources available, but their use always involves opportunity costs. For example, computing more to improve one partial plan may save a flight and its associated fuel and pilot costs, while computing more to improve another partial plan may save more lives. But one cannot simply just decide what computations to make by comparing the non-computational goods at issue because some computations (e.g., creating a fast local database) can speed or improve other computations, thus representing non-computational goods only indirectly.

The usual approach taken to managing these competitions is bureaucratic, with some central or higher authorities attempting to make broad allocations of resources and requiring the different organizations and their components to operate within these allocations and to explicitly request reallocations if such operation becomes impossible. Unfortunately, bureaucratic management of the resources for complex activities is too inflexible to meet the ordinary flow of events, which continually undermines the assumptions underlying such allocations by changing the demands on component organizations in ways difficult to predict by central planners. In particular, the most important competition for resources is not a static competition among relatively permanent organizations and their components: instead, the most important competition for resources is among an ever-changing set of tasks and portions of the overall activity. This competition arises as changing circumstances undermines portions of the plans of the organizations involved in the activity. To deal with these changing circumstances, the organizations and their components must continually repair the plans directing their current activities. We say repair rather than replace since in many cases traditional replanning from scratch may take too long to permit effective operation of the organizational components. The result is that the most important competition for resources is the competition among portions of the overall plan for the attentions of the authorities engaged in the activity. This competition is dynamic, cutting across organizational lines and changing much more rapidly than the structures of the organizations themselves (even when the organizations dynamically form teams to deal with problems). Bureaucratic management of resources does not appear adequate to deal with these dynamic resource competitions.

To manage large-scale activities and their dynamic resource competitions effectively, we must seek to address the competitions among tasks directly rather than indirectly through a fixed matrix of organizations. Instead of focussing on a fixed array of organizations and their capabilities, we must focus on the changing set of tasks, the importance of these tasks to each other (as accomplishing one task may facilitate or impede another), and the relevance of the organizational capabilities to each task. The question then becomes how to allocate resources to tasks, with the allocation of resources to organizations derived from the primary task-centered allocation.

2 Market allocation of heterogeneous resources

The field of economics has a well-known answer to the problem of allocating resources to tasks: use markets. It suggests viewing the competition for resources in terms of resource-endowed consumers that represent tasks or goals of the activity and resource-transforming producers that represent computational or reasoning procedures, informational resources, or noncomputational agents. Equilibrium ("market") prices correspond to allocations of resources to tasks that balance supply and demand, thereby ranking particular resource demands by their importance relative to other needs and the activity's ability to supply or produce the resources.

Economics bases its market prescription on practical, theoretical, and organizational grounds. In practice, markets for many goods provide extremely rapid response to changing circumstances (as observation of commodity markets shows) as well as flexible and measured responses to changing demands (which is not to say that overreactions are not possible). For a sizable range of important goods, markets empirically provide the best known method of allocation, even to the point of springing up and supplanting more bureaucratic mechanisms when not effectively suppressed. In addition, economic theory proves markets to be the most efficient allocation mechanisms possible, in the sense that the allocations determined by markets best satisfy the preferences of the agents involved. And markets do this while making smaller demands on the organizational structure of the activities using the resources, as they depend mainly on the calculations of the individual agents or organizational components competing for resources rather than on any calculation by central planning authorities.

Markets thus seem highly suited to complex resource allocation problems, as they do not require synoptic computational abilities and still achieve the best possible results. But their use in organizing activities has been limited for several reasons, especially the difficulty of recognizing competitions in a timely fashion and of tracking the wide variety of material and non-material goods involved in some activities. That is, markets in diesel fuel persist and succeed because diesel fuel is a commodity used by many agents over long periods of time, has very slowly changing or highly predictable properties, and is arbitrarily divisible, with different batches distinguished only by their size, not by other qualities. In contrast, it is difficult, at least within traditional organizational structures, to recognize the need for a market in strategies to repair a transportation plan upon the closure of an airfield [4], much less to identify the organizational agents that might appropriately enter into such a market. In consequence, traditional approaches to organizing large scale activities mix market and other mechanisms. The activities employ markets for resources when those markets already exist and are conveniently accessed by the component organizations, but use bureaucratic or other non-market mechanisms for the structuring the larger part of the activity. Traditional approaches thus are denied the benefits of market allocation for many parts of their activities.

3 Market-guided planning and search

Our work aims to change this situation dramatically by exploiting two developments in artificial intelligence: namely the development of computational markets and the development of automatic planning systems.

Limited forms of computational markets go back a long way in computer science; one can even view some of the early job-scheduling procedures for batch processing as implementing a "market" in computation for which users "bid" by means of commands on their job-control cards (punched cards, that is). But the true flowering of computational markets has occurred only recently, mainly with the WALRAS computational economy developed by Wellman [38]. Designed to make use of the main ideas of theoretical economics about markets, WALRAS provides a general mechanism for implementing markets in arbitrary sets of goods, traded by arbitrary consuming and producing agents.

WALRAS thus provides a mechanism for implementing markets in goods as they appear and disappear, as long as one can identify the need for such markets and the participants in them. Automatic planning systems supply means for making these identifications. Typical automatic planners take a set of tasks or goals and expand or refine the set by identifying methods for achieving some of these tasks or goals [9]. Now in fact, most planning systems do not stop there, but also search through different choices of which methods to use to achieve which goals, until a complete set of goals and corresponding methods is obtained that provides a comprehensive plan for achieving the initial goals. Unfortunately, when used in this way, automatic planning systems tend to replicate the traditional problems with bureaucratic organizations, as they produce bureaucratic solutions to problems: comprehensive plans clearly resemble traditional bureaucratic organizational structures, which present a fixed set of departments or positions addressing fixed tasks.

Our work aims to obtain the advantages of market mechanisms in a much wider range of situations than is traditional by making use of the mechanisms underlying automatic planning systems in a new way: not as means for constructing comprehensive overall plans, but as ways of identifying and constructing (or abandoning) portions of a market in tasks and the resources available for addressing them. That is, rather than have a planner itself choose the method to be used to achieve some goal, we use the planning mechanisms to create a market for that purpose, creating both the good of achieving the goal, identifying the producers of this good, i.e., the methods for achieving the goal, as well as the consumers of the good, i.e., the methods for achieving other goals which might rely or exploit the achievement of this one. The search for which assignments of methods best achieve the overall set of goals is then carried out by the market mechanism, which assesses the overall supply and demand for the market goods.

Where traditional planners make good choices of methods on a local basis, a market-guided planning system makes choices based on global assessments of importance, overcoming, we believe, the limitations of the traditional local choice methods. Improving the choices of traditional planners has usually been viewed as a problem of bringing more knowledge to bear on the individual choices of methods for achieving each task [6], and this has proved difficult as increasing complexity of the activity requires taking a larger and larger perspective in these choices. The advantage of the market-guided approach is that it provides a means for constructing comprehensive perspectives quickly using only application of the knowledge local to each method, with iterative adjustment of prices serving to combine and reconcile the local perspectives into a global one. Thus in the market-guided approach, no agent need know everything; the real knowledge is in the local methods, and the global market mechanisms serve only to combine the results of this knowledge in summary form. Other approaches, such as neural networks, share something of the flavor of the market-guided approach, but unless constructed to mirror the structure of a market (which none seem to have been), lack the assurances of efficiency and rationality provided by the economic theory of markets.

4 RECON, the reasoning economy

We determine rational allocations of the full range of resources by using an artificial market economy implemented as RECON, the Reasoning ECONomy, to determine prices or trading ratios among resources. Since computations like plan construction or revision may change the resource demands of the activity, the relevant competitions for resources involve computational resources (time, database access or update control, etc.) as well as the more obvious non-computational resources (fuel, aircraft, etc.).

Our preliminary implementation of RECON is built in CLOS on top of Wellman's WALRAS artificial economy [38]. WALRAS provides the basic market notions of consumer and producer agents, goods, and auctions for these goods. Each consumer and producer enters into a subset of the auctions. Consumers are endowed (possess) bundles of various goods, and enter into the auctions for each of the goods in their endowment. Producers do not possess goods, but instead transform a set of input goods into an output good, and enter into the auctions for their input and output goods. Consumers and producers place bids in each of the auctions in which they participate. These bids may be as simple as indicating the desire to trade a specific quantity of the good at the current price; or as complex as a full schedule of amounts to trade (buying or selling) at each possible price. Consumers and producers determine their bids in different ways. Consumers have a utility function or preference order over possible bundles of goods, and bid so as to trade a less preferred bundle for a more preferred bundle. Producers have a production function that describes how much output derives from each bundle of inputs, and use this production function to determine bids for input goods from prices for output goods or vice versa. WALRAS determines equilibrium prices by an iterative procedure that at each iteration determines the total supply and demand for each good at current prices and adjusts the prices accordingly if supply and demand do not match. In observations, WALRAS computes equilibrium or near-equilibrium prices in only a few iterations. WALRAS makes a number of special economic assumptions about agents, the most important one being that agents are "price taking", that is, no agent is large enough in the market to influence prices simply by its own actions. This assumption is actually not true in many of the small reasoning markets of interest, but its failure has not yet been shown to matter in practice due to the iterative nature of equilibriation.

As is apparent from this description, the economic structures provide by WALRAS are not specific to any types of goods, consumers, or producers. RECON builds on WALRAS by augmenting these generic notions with notions of good, consumer, and producer specific to reasoning tasks, and by introducing computation goods to allocate to different tasks. Its main additions to WALRAS are a good representing computation, the notions of computation consumers and producers, a taxonomy of action types, and an execution mechanism that takes actions in an order determined by bids for computation. The preliminary RECON implementation schedules tasks on a single processor by auctioning off opportunities to compute. Computation consumers and producers bid for amounts of the computation good as though they were bidding for different amounts of time. Computation consumers have a computation function in addition to their utility function, where their computation function indicates their use of computation in different circumstances. Computation producers take computation as input and produce outputs, with different subtypes of computation producers characterized by their computation production functions that indicate how much computation they take to produce their results (i.e., their running time). The taxonomy of action types introduces the distinction between actions and states, and differentiates a standard variety of action types reflecting mainly planning and reason maintenance actions. Base level actions have associated procedures, and the execution mechanism schedules the execution of these procedures by using the bids each procedure's producer bids for computation. However, though the bids reflect the amounts of computation desired by consumers and producers, the execution mechanism in fact provides uninterrupted opportunities to compute as long as desired (possibly forever) rather than the opportunity to compute only as much as the bid specified. The selection of the producer to execute is made by one of several methods, the simplest being to simply pick the producer bidding for the maximum amount of computation among all producers bidding for computation. This corresponds intuitively to taking actions to make as much progress as possible at each step. The implementation provides a fairly manual method for charging producers and the consumers they serve for the computation consumed once a producer's procedure is executed, and a fairly manual method of "garbage collection" for removing bankrupt or otherwise useless consumers and producers from the economy when these charges remove all the endowment of consumers.

5 The market-guided reason maintenance system

We developed and tested the preliminary RECON architecture in conjunction with a market reason maintenance system, MRMS. MRMS reorganizes the operations of a traditional RMS [10] into local revisions that are done only upon determination of need, rather than the automatic and global revisions performed by traditional reason maintenance systems [14]. To observe the system's behavior and perform preliminary tests of the implementation, we took the AMORD pattern-directed procedure system [6] and replaced its traditional RMS with MRMS. We also replaced AMORD's first-in, first-out queue of procedure invocations with the RECON producer-scheduling mechanism, and redescribed AMORD goals and methods in terms of computation consumers and producers. The AMORD goals thus corresponded to computation consumers that were endowed with computation and would trade computation for achievement of the goal, while AMORD methods corresponded to computation producers that would peform procedures or draw conclusions given the computation to do so.

The preliminary RECON implementation has many limitations discussed below, and AMORD provides only a very simple platform for conducting and studying reasoning, but even this primitive experimental apparatus was sufficient to observe some of the desired phenomena of reasoning economies. For example, depending on the amount of computation required by the basic reason maintenance procedures relative to each other, one could observe a series of reasoned updates either being performed sequentially as individual revisions (which take little time but may leave the network of justifications somewhat incoherent), or being performed indirectly by avoiding any revision during the incoming series but revising the entire set at once in a larger update operation at the end (which can take more time but leaves the network of justifications mutually coherent). This illustrates one of the strenths of the market-oriented approach: tasks which have implications for many other tasks are automatically accorded great importance, in some cases even more than tasks that are very important on their own. (An analogous example might be that using a vehicle to transport supplies of disinfectant to a hospital might be more important in some circumstances than using the vehicle to transport wounded, even though one naturally would view a wounded person as more important than ordinary hospital supplies.) These determinations derive purely from the information local to each task and method and from the pattern of interaction of the tasks and methods as reflected in the market topology. There is no need to apply specific knowledge to recognize the importance; the market mechanisms provide that indication automatically.

Our experience with the RECON/MRMS/AMORD combination point out the need for several improvements and extensions to RECON. The most urgent improvements concern the transaction mechanism for auctioning off computation. The preliminary implementation allocates uninterruptible computation opportunities, but for flexible coordination of activities, what is desired is instead allocation of more definite segments of time (e.g., uniform time slices) with differing degrees of interruptibility (uninterruptible, interrupt with suspension of state, abort and restart, anytime procedures, etc.). Our student Nathaniel Bogan is incorporating several standard approaches from economic theory and practice into RECON [3]. This literature offers important results applicable to computation and other special sorts of resources. (Unlike fuel oil, for which one has many suppliers and consumers and many pieces which can be traded and consumed independently, time on a processor is unique; only one process may compute at any given moment.) Bogan uses these results to design auction mechanisms appropriate to the special properties of computation and other goods.

In addition, the preliminary implementation of RECON provides only an ad-hoc mechanism for dealing with bankruptcy. For the sort of expectation-guided allocation we envision, we require more sophisticated means of dealing with tasks that overrun their allocation and thus consume more resources than they had available to trade. Bogan is also replacing these ad-hoc mechanisms with more reasonable approaches to bankruptcy, ranging from automatic removal of the bankrupt consumers and any associate special-purpose producers from the reasoning economy, to the addition of standard mechanisms for suspending and resuming execution of procedures as more resources become available.

Another limitation of the preliminary implementation is that it allocates only one computational good on a single processor. To perform the multi-good, multi-agent coordination needed for large-scale activities, we plan to expand the notion of computational good to include space and other resources on multiple servers.

A final limitation is that the current taxonomy of action types does not cover the full range of typical operations of reasoning, planning, and search procedures. We are expanding and refining this taxonomy as the need arises. This expansion more generally provides strong motivations for developing a suitable language for describing computational agents and their economic properties.

5.1 Representing preference information

Preferences constitute one of the basic elements of information about economic agents, and a key problem in both our investigation of market-guided planning and in decision-theoretic planning in general is finding a good representation of preference information. Toward this end, we have developed a qualitative logic of preference *ceteris paribus* or preference "other things equal" [13, 15, 39] in collaboration with Michael Wellman.

The most common representation used or assumed for preference information is that of numerical utility functions. These have the advantage of offering complete comparisions of the relative desirability of all alternatives and of applying numerical optimization procedures in decision making. But they also suffer from severe problems of convenience and generality, especially when attempting to integrate them into the automatic reasoning and planning system developed in the artificial intelligence literature. In the first place, numeric utility functions are too specific: the foundations of decision theory and economics start with qualitative preference orders, and many different utility functions can represent the same preference order [30]. Thus a utility function may convey more information than is really there. In the second place, while utility functions make some optimization procedures convenient, they tend to make specification and elucidation of preferences difficult. Most people exhibit a well-known aversion to expressing their judgments in numerical terms, yet are often willing to provide qualitative expressions of preference with confidence. Working directly with utility functions thus makes people hesitant to supply the necessary information, and also provides no way to make use of the qualitative comparisions that informants might provide immediately. In the third place, utility functions provide no connection with the notion of *goal* upon which most reasoning and planning systems in use in artificial intelligence are based. Most of these systems solve problems and construct plans in response to stipulations of one or more propositions representing the goals to be achieved. Decision theory has no notion of "things to be achieved", only rankings of the relative desirability of outcomes, while goal-based systems frequently have only the notion of "things to be achieved" and no way of comparing the relative desirability of potential solutions or plans.

In response to these problems with utility functions, we have been working toward finding ways of representing information about preferences that permits expression of both standard decision-theoretic information as well as standard goal-based specifications, so as to be able to exploit the strengths of both decision theory and artificial intelligence. More generally, we have aimed to find preference representations capable of encoding and using whatever preference information is available, without having to wait until enough is specified to determine a utility function.

Our main results so far have been very encouraging. Our qualitative logic of preference *ceteris paribus* provides a uniform language in which one can express both ordinary decision-theoretic preferences as well as the standard notion of goal, which we interpret to mean conditions preferred to their opposites other things equal. (It is easy to show that one cannot reasonably interpret goals as conditions preferred to their opposites without qualification, as that interpretation trivializes planning with multiple independent goals. In particular, that interpretation forces all states that satisfy some goals but not others to be indifferent; so that one cannot add any information to the goals to state that, for example, outcomes satisfying 5 goals are preferable to outcomes satisfying only 4.) We are continuing to develop the theoretical structure and inferential capabilities of this logic and some close variants, but the basic language already provides a useful tool for encoding qualitative preferential information. Furthermore, it appears that many important decisions can be made simply on the basis of dominance arguments expressed completely qualitatively, so this approach should permit important decision-making and planning to proceed even without numeric utilities. Wellman [37] has developed the probabilistic version of this idea into a dominance-guided planning procedure that plans "up to tradeoffs" in his terminology, meaning it constructs plans that are optimal as far as the qualitative probabilistic information is concerned, and we expect an even more robust planning procedure could be constructed in a similar fashion using only qualitative preference information.

Of course, a rich language for encoding preference information would include quantitative representations as well, and we are working toward a preference language that spans the spectrum from completely qualitative representations like our language of comparative preference to ordinary numeric utility functions, including intermediate representations of multiattribute utility functions such as subutility composition trees [40], the standard forms of multiattribute utility functions [25], and their application to expressing different types of planning goals [18, 19].

6 Related work

Work on distributed AI has made use of market notions for about a decade, but these uses have been more suggestive than substantive until very recently [26]. The early notion of contract net of Davis and Smith [5], for example, appealed to market metaphors (i.e., bidding for contracts) without providing means for using prices in its protocols, and the "society of mind" theory of Minsky [27] takes a broad but informal view of mental structure and activity that subsumes market activity and many other forms of interaction. The more substantive recent efforts of Waldspurger and others [35] have focussed on allocating single goods, namely processor time in a distributed operating system. Wellman's WALRAS system [38], upon which our current implementation is based, constitutes the first general and substantial use of market notions, providing a computational mechanism for finding overall equilibria of multi-good, multi-agent markets. To date, Wellman has primarily studied the theoretical properties of the specific WALRAS mechanisms and the use of WALRAS to allocate non-computational goods.

Work on planning and replanning has traditionally focussed on centralized planning methods that do not scale well to construction and maintenance of plans for large-scale activities. Most of this work treats planning only as a problem of achieving discrete "goals", with little or no information available to compare the value of achieving one set of goals with the value of achieving another. Consequently, the primary means used to guide the planning and replanning process have been purely structural (how many goals achieved) or probabilistic (how likely the goals are achieved). Recent efforts on decision-theoretic planning and acting [8, 9, 11, 19, 31, 37] have attempted to employ utility information, but most planning methods in use remain fairly insensitive to this information. Even in those efforts, the focus is on optimizing the non-computational properties of the plan without regard to the computational resources required to do so.

Work by Russell and Wefald [28, 29] and others [1, 7, 20, 21, 22, 23, 24] on rational allocation of computational effort in search has complemented this work on planning by studying optimal allocation of computation time in search, primarily in the simpler domain of game-playing search. This work presumes a utility model for non-search actions, and seeks to trade off search time and non-search utility. The specific mechanisms investigated to date, however, do not distinguish well the variety of computational resources and goals of interest. Our students Ronald Bodkin and Michael Frank investigated several ideas in rational guidance of search, developing some new search methods [2] and exposing severe problems with the assumptions underlying extant rational search procedures [16, 17].

While work on qualitative logics of preference goes back some time in philosophical logic [32, 33], work on the problem in artificial intelligence seems to originate with our contributions [39, 13]. However, this work has its roots in Wellman's earlier work on decomposition of multiattribute utility functions [36], which we also further developed in [40]. Work on a related problem, that of constructing multiattribute utility functions that model important types of planning goals, has also been under investigation by Haddawy and Hanks [18, 19].

Most treatments of reason maintenance have followed the original [10] in performing complete updates, and theories of reasoned revisions have followed in making this sometimes impractical assumption. Our mathematical work on reasoned assumptions [12] provides a better understanding the properties of MRMS and more general reason maintenance systems that perform piecemeal revisions of their information. In particular, it develops the strong connection between reasons and preference information to show how one may view reason maintenance systems as performing rational reasoning.

Acknowledgments

We thank Michael Wellman for making his WALRAS system available to us for use in this project, and for his consultation in our reorganization of its code. This research was supported by ARPA and Rome Laboratory under contract F30602-91-C-0018.

References

[1] E. B. Baum and W. D. Smith. Best play for imperfect players and game tree search. NEC Research Institute, May 1993.

[2] R. J. Bodkin. Extending computational game theory: Simultaneity, multiple agents, chance and metareasoning. Master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, Sept. 1992.

[3] N. R. Bogan. Economic allocation of computation time with computation markets. Master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1993. Thesis Proposal.

[4] S. Cross. A proposed initiative in crisis action planning. DARPA Information Science and Technology Office, May 1990.

[5] R. Davis and R. G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20:63–109, 1983.

[6] J. de Kleer, J. Doyle, G. L. Steele Jr., and G. J. Sussman. AMORD: Explicit control of reasoning. In *Proceedings of the ACM Symposium on Artificial Intelligence and Programming Languages*, pages 116–125, 1977.

[7] T. Dean. Decision-theoretic control of inference for time-critical applications. *International Journal of Intelligent Systems*, 6(4):417–441, 1991.

[8] T. Dean and M. Boddy. An analysis of time-dependent planning. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 49–54, 1988.

[9] T. L. Dean and M. P. Wellman. *Planning and Control*. Morgan Kaufmann, San Mateo, CA, 1991.

[10] J. Doyle. A truth maintenance system. *Artificial Intelligence*, 12(2):231–272, 1979.

[11] J. Doyle. Artificial intelligence and rational self-government. Technical Report CS-88-124, Carnegie-Mellon University Computer Science Department, 1988.

[12] J. Doyle. Reasoned assumptions and rational psychology. *Fundamenta Informaticae*, 20, 1994.

[13] J. Doyle, Y. Shoham, and M. P. Wellman. A logic of relative desire (preliminary report). In Z. W. Ras and M. Zemankova, editors, *Methodologies for Intelligent Systems, 6*, volume 542 of *Lecture Notes in Artificial Intelligence*, pages 16–31, Berlin, Oct. 1991. Springer-Verlag.

[14] J. Doyle and M. P. Wellman. Rational distributed reason maintenance for planning and replanning of large-scale activities. In K. Sycara, editor, *Proceedings of the DARPA Workshop on Planning and Scheduling*, pages 28–36, San Mateo, CA, Nov. 1990. Morgan Kaufmann.

[15] J. Doyle and M. P. Wellman. Representing preferences as *ceteris paribus* comparatives. In S. Hanks, S. Russell, and M. P. Wellman, editors, *Proceedings of the AAAI Spring Symposium on Decision-Theoretic Planning*, 1994.

[16] M. P. Frank. Rational abstract search. MIT Lab for Computer Science, Nov. 1993.

[17] M. P. Frank. Rational partial evaluation of decision trees. Master's thesis in progress, 1994.

[18] P. Haddawy and S. Hanks. Issues in decision-theoretic planning: Symbolic goals and numeric utilities. In *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control*, pages 48–58, 1990.

[19] P. Haddawy and S. Hanks. Representations for decision-theoretic planning: Utility functions for deadline goals. In B. Nebel, C. Rich, and W. Swartout, editors, *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning*, pages 71–82, 1992.

[20] O. Hansson and A. Mayer. The optimality of satisficing solutions. In *Proceedings of the Fourth Workshop on Uncertainty in Artificial Intelligence*, 1988.

[21] O. Hansson and A. Mayer. Decision-theoretic control of search in BPS. In *Proceedings of the Symposium on AI Limited Rationality*, pages 59–63. AAAI, 1989.

[22] O. Hansson and A. Mayer. Heuristic search as evidential reasoning. In *Proceedings of the Fifth Workshop on Uncertainty in Artificial Intelligence*, pages 152–161, 1989.

[23] E. J. Horvitz. Reasoning under varying and uncertain resource constraints. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 111–116, San Mateo, CA, 1988. AAAI, Morgan Kaufmann.

[24] E. J. Horvitz, G. F. Cooper, and D. E. Heckerman. Reflection and action under scarce resources: Theoretical principles and empirical study. In N. S. Sridharan, editor, *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, volume 2, pages 1121–1127, San Mateo, CA, August 1989. International Joint Conferences on Artificial Intelligence, Inc., Morgan Kaufmann.

[25] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley and Sons, New York, 1976.

[26] M. S. Miller and K. E. Drexler. Markets and computation: Agoric open systems. In B. A. Huberman, editor, *The Ecology of Computation*, pages 133–176. North-Holland, Amsterdam, 1988.

[27] M. Minsky. *The Society of Mind*. Simon and Schuster, New York, 1986.

[28] S. Russell and E. Wefald. Principles of metareasoning. *Artificial Intelligence*, 49(1-3):361–395, May 1991.

[29] S. J. Russell. *Do the Right Thing: Studies in Limited Rationality*. MIT Press, Cambridge, MA, 1991.

[30] L. J. Savage. *The Foundations of Statistics*. Dover Publications, New York, second edition, 1972.

[31] D. E. Smith. A decision theoretic approach to the control of planning search. Technical Report LOGIC-87-11, Department of Computer Science, Stanford University, 1988.

[32] G. H. von Wright. *The Logic of Preference: An Essay*. Edinburgh University Press, Edinburgh, 1963.

[33] G. H. von Wright. The logic of preference reconsidered. *Theory and Decision*, 3:140–167, 1972. Reprinted in [34].

[34] G. H. von Wright. *Philosophical Logic: Philosophical Papers, Volume II*. Cornell University Press, Ithaca, NY, 1984.

[35] C. A. Waldspurger, T. Hogg, B. A. Huberman, J. O. Kephart, and S. Stornetta. Spawn: a distributed computational economy. *IEEE Transactions on Software Engineering*, 1992.

[36] M. P. Wellman. Reasoning about preference models. TR 340, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, May 1985.

[37] M. P. Wellman. *Formulation of Tradeoffs in Planning Under Uncertainty*. Pitman and Morgan Kaufmann, 1990.

[38] M. P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.

[39] M. P. Wellman and J. Doyle. Preferential semantics for goals. In *Proceedings of the National Conference on Artificial Intelligence*, pages 698–703, 1991.

[40] M. P. Wellman and J. Doyle. Modular utility representation for decision-theoretic planning. In *Proceedings of the First International Conference on AI Planning Systems*, 1992.