# Formalizing Reasoning About Change: A Qualitative Reasoning Approach
## (Preliminary Report)

**James M. Crawford** **David W. Etherington**

AI Principles Research Department
AT&T Bell Laboratories
600 Mountain Ave.
Murray Hill, NJ 07974-0636
{jc,ether}@research.att.com

## Abstract

The development of a formal logic for reasoning about change has proven to be surprisingly difficult. Furthermore, the logics that have been developed have found surprisingly little application in those fields, such as Qualitative Reasoning, that are concerned with building programs that emulate human common-sense reasoning about change. In this paper, we argue that a basic tenet of qualitative reasoning practice—the separation of modeling and simulation—obviates many of the difficulties faced by previous attempts to formalize reasoning about change. Our analysis helps explain why the QR community has been nonplussed by some of the problems studied in the nonmonotonic reasoning community. Further, the formalism we present provides both the beginnings of a formal foundation for qualitative reasoning, and a framework in which to study a number of open problems in qualitative reasoning.

## Introduction

Formalizing reasoning about change has received much attention in the nonmonotonic reasoning community (NMR) [Ginsberg 1987]. The qualitative reasoning community (QR) has focused on the closely-related task of developing programs that emulate human common-sense reasoning about physical systems [Weld & de Kleer 1990]. Strangely, there has been little interaction between the two fields. This may be partially explained by the gulf between the formal methods used in NMR and the experimental methods used in QR. However, certain principles have recently begun to emerge in QR that are useful in guiding the formalization of reasoning about change. In particular, it has become apparent that reasoning about change can be simplified if one separates the task of (formally) deriving a description of a particular system from general principles (the "modeling problem"),[1] and the problem of making predictions based on that description (the "simulation problem"). We describe a preliminary formalism incorporating this separation. We show that this formalism improves on previous work by showing that it correctly handles examples in which information is available about non-initial states (*e.g.*, the "stolen car problem" [Kautz 1986] and its relatives).

Given a complete description of a system (expressed, for example, as a differential equation or a set of statements in first-order logic), it is relatively straightforward to predict the possible behaviors of the system (using a numerical simulator, a program such as QSIM [Kuipers 1986], or a formalism such as that presented in [Morgenstern & Stein 1988]). Such predictions, in general, take the form of state transition graphs that show the effects of each possible action in each possible state. Given the initial state of the system, one can then use the state transition graph to predict the behavior of the system. We refer to the task of deriving state transition graphs (and making behavioral predictions) as the *simulation* problem.

Things are more complex if one does not have a complete description of the system. A description may be incomplete, for example, because one wants to allow for the possibility that events may fail to have their expected effects (*e.g.*, flipping the switch may fail to light the lamp) or because one wants to allow for the possibility of 'miracles'—changes in the state of the

[1] Unfortunately, the term *model* has very different meanings in QR and NMR. In QR, a model is a description of a component or system (*e.g.*, the normal model of a light is that you turn it on and it glows, while the fault model is that you turn it on and nothing happens). In NMR, a model is an assignment of truth values. The sense of model we mean will (we hope) be clear from context (*e.g.*, whenever we speak of "model building" we are using model in the QR sense). When the sense is not obvious we use "qualitative model" or "logical model".

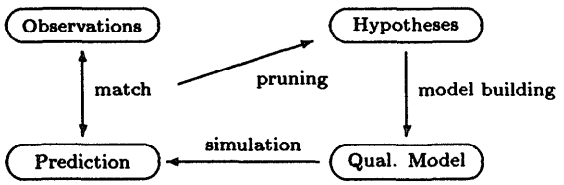


Figure 1: The Roles of Model Building and Simulation.

world that are not caused by any known actions. In such cases, one must hypothesize whether each event has its normal effects and which—if any—miracles occur at each time. Any such set of hypotheses induces a complete hypothetical description (a 'model') of the system. We refer to the task of choosing a set of hypotheses as ***model building***. Each model can be simulated to derive behavioral predictions, which can then be compared against observations of the actual system. Sets of hypotheses that induce models whose predicted behavior does not match the observed behavior of the system can then be pruned away (see Figure 1).[2]

The nonmonotonic reasoning community has generally failed to distinguish the model building and simulation problems. As a result formalisms have been developed that allow observations to corrupt the simulation machinery. This leads to difficulties when information about non-initial states is available: rather than allowing models to be pruned because they do not match observations, existing formalisms (*e.g.*, [Baker 1989]) essentially change the behavior of the "simulator" so that any model is made to match the observations. An example of this problem is given after the next section.

By explicitly separating the model-building task from the simulation task, we gain two advantages. First, we are able to handle a class of problems that has caused difficulties for previous approaches. Second, we lay the groundwork for a formalization of qualitative reasoning practice. We are thus able to understand why problems such as the Yale Shooting Problem have not been of interest in QR, and suggest future work that could be of interest to both communities.

## Nonmonotonic Formalisms for Reasoning About Change

In this section we briefly review the ideas from past work on formalizing change that are necessary to understand the discussion that follows.

Much formal work on representing change is based on the situation-calculus [McCarthy & Hayes 1969]. The situation calculus represents the world in terms of situations (states of the world), actions, and fluents (time-varying properties that may or may not hold in particular situations). The predicate *Holds* is a relation between situations and fluents that states that a fluent holds in a particular situation. The function *result* maps a situation and an action to the situation produced by performing the action in that situation. Axioms are usually given that detail the necessary preconditions for, and results of, the actions.

Logical formalisms for actions are generally nonmonotonic for two reasons. First, one may want to state that an action *typically* has certain effects, but be unable or unwilling to explicitly list all combinations of circumstances in which the action may fail to have these effects. Second, actions often have indirect effects, so it is in general infeasible to axiomatize the exact effects of successful actions. (For example, painting a house probably changes the value of its color fluent, but it may also change the value of the fluent *pretty*). To address this problem, *persistence axioms* are usually added that state that if a fluent value is not known to change as the result of an action, then one may assume that it persists. Both types of nonmonotonicity are generally represented by the use of *abnormality* predicates. The extensions of these predicates are then minimized; logical models with fewer abnormalities (*i.e.*, fewer violations of default assumptions) are preferred over those with more abnormalities.

Thus, for example, a (simplified) axiom stating *if block y is initially clear, then block x will be on it after the action PutOn(x, y)* might look like

$$\begin{aligned}&\forall s.\ \forall xy.\ Holds(clear(y), s)\\ &\qquad \wedge \neg Abnormal(s, PutOn(x,y))\\ &\qquad \supset Holds(On(x,y), result(PutOn(x,y), s))\ .\end{aligned}$$

This axiom says that if $y$ is clear, and the *PutOn* proceeds normally, then $x$ will be on $y$ in the resulting state.

Of course, there is the potential for conflict between the statement that fluents normally persist and the statement that actions normally change certain fluents.[3] We shall return to this conflict in the next section.

## A Problematic Example

Consider the example shown in figure 2. Initially there is water in $Tank_1$. After waiting for some period of time the valve is opened. We then ask whether there is water in $Tank_2$.[4] We can represent this problem using the fluents $Full(Tank_1)$ and $Full(Tank_2)$, and the actions *Wait* and *Open* (*i.e.*, open the valve). We

[2]The idea of simulating all possible models and then pruning those that conflict with observations is the heart of the QR system MIMIC [Dvorak and Kuipers, 1989].

[3]In fact, the question of how to resolve conflicts between defaults is one of the fundamental questions in the study of nonmonotonic reasoning.

[4]The astute reader may notice that this example is similar to an unnameable example often studied in NMR.

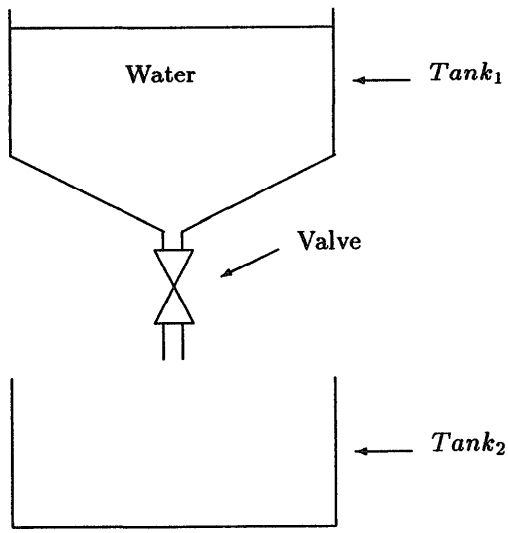


Figure 2: The tank problem.

axiomatize *Open* as

$$\forall s.\ Holds(Full(Tank_1), s) \wedge \neg Abnormal(s, Open) \supset Holds(Full(Tank_2), result(Open, s))$$

thus allowing for the possibility that the valve may fail to open. Assume for the moment that we do not separate model building from simulation (*e.g.*, that we use Baker's formalism [Baker 1989], which handles many of the problems studied in non-monotonic reasoning, including the original Yale Shooting Problem). Note that initially we have $Full(Tank_1)$ and $\neg Full(Tank_2)$. The persistence assumption is violated if $Full(Tank_2)$ becomes true after *Open*, but it violates the default that the valve usually works for $Full(Tank_2)$ to remain false. Although one or the other violation must occur, a straightforward axiomatization gives no basis for preferring one to the other. Such ambiguities will occur whenever the effects of actions are defeasible.

The usual solution to such difficulties is to minimize persistence abnormalities with lower priority than other abnormalities (in effect preferring persistence failures to action failures). This allows *Open* to be normal and thus to change the value of $Full(Tank_2)$. Unfortunately such a prioritization leads to other difficulties. For example, consider what happens if we know that $Tank_2$ remains empty after the valve is opened. The prioritized approach then allows us to conclude (unambiguously) that $Tank_1$ became empty during the *Wait* action, since that involves only a violation of a persistence assumption (in particular the persistence of $Full(Tank_1)$). Now, while this is *possible*, it is hardly the only reasonable conclusion (especially since we explicitly axiomatized the possibility that valves can fail). In general, such a prioritization can force *any number* of persistence assumptions to be dropped in order to avoid introducing a single non-persistence abnormality—the simulation is changed to make the model "work"!

Things are better if we separately simulate the case in which the valve works, and the case in which it fails, and then check the resultant predictions against observations. If we hypothesize that the valve works, we can simulate the expected behavior and predict that $Tank_2$ fills. Similarly, if the valve fails then $Tank_2$ remains empty. In either case, water remains in $Tank_1$ during the *Wait*. We can then compare these predictions against the observation that $Tank_2$ is empty, and rule out the case in which the valve works. Note that the key difference between this approach and the prioritized approach is that we explicitly separate consideration of each possible set of assumptions about the normality of the actions (which we will refer to as *modeling assumptions*) from the simulation process (which determines persistence abnormalities). In the next section we show how such a distinction can be formalized; later, we return to the tank problem and show how it is solved by our approach.

## Formalizing the Modeling/Simulation Distinction

In this section, we outline a formalism for reasoning about change that explicitly separates modeling and simulation. We describe the axiomatization, discuss the distinguished role of *observations*, and explain how inference is performed.

For concreteness and simplicity, we base our formalism on the situation calculus. However, the basic idea of separating modeling and simulation is a much more general notion, and does not depend on the situation-calculus (*e.g.*, QPE [Forbus 1984] and QPC [Crawford, Farquhar & Kuipers 1990] both make such a separation, but allow reasoning about non-instantaneous and overlapping processes, and describe physical devices in terms of continuously-varying parameters).

### Axiomatization

**Actions:** We use $MAb(action, time, fluent)$ to mean that the instance of action *action* at time *time* is subject to a modeling abnormality that affects fluent *fluent*. Thus, if $Preconditions(Action, sitn)$ stands for a formula that expresses the preconditions for action *Action* in situation *sitn*, then

$$\forall sitn.\ Preconditions(Action, sitn) \wedge \neg MAb(Action, Time(sitn), Fluent) \supset Holds(Fluent, result(Action, sitn))$$

expresses the fact that the normal model of the action is one where *Fluent* holds after the action (if the preconditions hold before). The particular form of the axiom is relatively unimportant and can be changed

to reflect the desired behavior; the key feature is the fact that modeling abnormalities are properties of the time at which a situation occurs, rather than of the situation itself.

**Persistence:** By contrast, we view persistence abnormalities as properties of situations (*i.e.*, as properties of the underlying fabric of the representation), and hence deal with them during simulation. Following Baker [1989] axioms are added that force the existence of situations corresponding to each possible set of fluent values.[5] We use $Ab(action, sitn, fluent)$ to mean that the performance of *action* in state *sitn* changes the value of *fluent* (this is "abnormal" since fluents are normally assumed to persist). Persistence is thus postulated by the usual single axiom:

$$\begin{aligned}&\forall action, sitn, fluent.\ Holds(fluent, sitn)\\ &\quad \wedge \neg Ab(action, sitn, fluent)\\ &\quad \Longleftrightarrow Holds(fluent, result(action, sitn))\end{aligned}$$

**Miracles:** In order to represent rich worlds with apparently uncaused changes (such as the Stolen Car Problem discussed in the next section), it is useful to have a representation for so-called "miracles"—changes that are not the result of any particular action (*e.g.*, the disappearance of one's car during some *Wait* event) [Lifschitz & Rabinov 1989]. Miracles are a type of modeling abnormality. They are treated in a slightly peculiar way:

$$\begin{aligned}&\forall sitn.\ Preconditions(Action, sitn)\\ &\quad \wedge Miracle(Action, Time(sitn), Fluent)\\ &\quad \supset Holds(Fluent, result(Action, sitn))\ .\end{aligned}$$

This axiom can be read as saying that, under certain circumstances (specifically, when *Preconditions* holds), in an abnormal model (one where a miracle occurs), *Fluent* will come to hold after the action *Action*.[6] Thus, while action axioms say that actions normally have certain effects, miracle axioms say that miracles *abnormally* have certain effects—that miracles are not the normal state of affairs. It should be noted that miracle axioms differ from typical axioms involving abnormalities in that *Miracle* appears positively in the antecedent. It might appear that there is thus no way to *infer* that a miracle has occurred. However, as we shall see when we discuss the stolen car example, models are pruned by observations, and in some cases all models without miracles may be rejected.

## Observations

An *observation* is any formula that mentions a specific (or an existentially-quantified) time or situation. Thus observations include facts about initial, final, and intermediate conditions, statements about what actions occur when, and even conditional statements about actions that behave differently at "landmark" times.

Observations have a distinguished role: they do not participate in simulation, but are critical in pruning the set of possible models. This role is described in the next sub-section.

## Inference

The actual reasoning process proceeds in three stages, which correspond to (1) simulating all possible qualitative models, (2) comparing the results with observations to rule out inappropriate models, and (3) choosing preferred models from among those remaining. These 3 stages are realized by

1. Minimizing *Ab* varying *result*, but holding *MAb* and *Miracle* fixed;
2. Augmenting the resulting theory with the observations; and
3. Minimizing *MAb* and *Miracle* in the augmented theory, varying *Ab* and *result*.

These stages are formally defined using circumscription in [Crawford & Etherington 1992].

Recall from our discussion of persistence axioms that we include Baker's axioms to force the existence of all possible situations. The interior (first) minimization essentially runs the simulation (specifying the *result* of each action in each situation, while minimizing persistence abnormalities) for each qualitative model. In this minimization, *MAb* and *Miracle* are fixed in order to force consideration of all qualitative models. Thus, for each configuration of modeling assumptions (*MAb*s and *Miracles*), a set of *Ab*-minimal logical models is produced. These logical models correspond to the results of simulation under those qualitative modeling assumptions.[7] The addition of observations then rules out configurations of modeling assumptions that, when simulated, yield predictions that conflict with actual observed values. Finally, *MAb* and *Miracle* are minimized to prefer those qualitative models that violate as few modeling assumptions as possible (while still satisfying the observations).[8]

In summary, the formalism simulates all possible qualitative models, and then prunes those that do not jibe with the observations. Among the remaining models, those that violate as few modeling assumptions as possible are preferred. The formalism differs from past

[5] [Crawford & Etherington 1992] details these axioms.

[6] The dependence of miracles on particular actions can, of course, be avoided by quantifying over actions.

[7] The result is similar to a total envisionment [Forbus 1984] in QR.

[8] *Ab* is varied in this minimization to prevent the persistence violations induced by simulation from unduly influencing the choice of models by making models incomparable; this avoids artifacts induced by differing qualitative models having different simulations. Allowing *Ab* to vary is not problematic since the functional relationship between modeling and simulation assumptions has already been determined in the first minimization.

work in that simulation (the minimization of $Ab$) is separated from model building (the minimization of $MAb$ and $Miracle$). Further, observations are included only after the minimization of $Ab$ (thus preventing them from affecting the simulation).

By taking a model-theoretic view, we can sharpen our understanding of this inference process. Models of the action and persistence axioms can be grouped into equivalence classes according to the sets of $MAb$s they contain. Within each equivalence class, models can be partially-ordered by their sets of $Ab$s. The first minimization then selects those models in each equivalence class that are $Ab$-minimal, and prunes the rest. The second step prunes models that are incompatible with the observations. Notice that this can result in entire equivalence classes being pruned (even though some members of the class—that are not $Ab$-minimal—may have been consistent with the observations). The final step selects, from the models that remain, those minimal in $MAb$.

Since each step begins with the models remaining after the previous step, and since the observations are not considered until after the $Ab$s for each configuration of $MAb$s are determined, it is easy to show that the introduction of observations cannot accommodate observations by generating abnormality in the simulation. This is what we set out to guarantee.

## Two Examples

### The Two Tank Problem

We now return to the two tank problem shown in Figure 2.[9] The first minimization results in two relevant $Ab$-minimal logical models,[10] one with no $MAb$s, in which $Open$ is effective (the valve works, and $Tank_2$ fills), and another with one $MAb$, in which the valve sticks (and $Tank_2$ remains empty). If there are no observations about the final state, both of these models are considered for minimization of $MAb$, and we conclude that $Tank_2$ fills. On the other hand, if we know that $Tank_2$ is empty after the $Wait$ and $Open$, the first model is pruned, but the second model remains: we explain the lack of water in $Tank_2$ by postulating that the valve stuck.

If one adds an axiom permitting the "miraculous" disappearance of water, there will be an additional $Ab$-minimal model, with one modeling abnormality, in which $Tank_2$ remains empty because $Tank_1$ emptied during the $Wait$. Given no observations about the final state, this model would be dropped in favor of the model with no modeling abnormalities. However, if one knows that $Tank_2$ remains empty, then the possibility of miraculous emptyings introduces an ambiguity between the model in which the water disappears and the model in which the valve fails. Notice that in neither of these models does the water disappear simply due to a failure of persistence. Something (a miracle in this case) has to *happen*. One might wish to further prioritize the minimization of modeling abnormalities to stipulate that miracles are much less likely than other types of modeling abnormalities; this presents no difficulties, and would cause one to attribute the lack of water in the $Tank_2$ to the failure of the valve rather than to the disappearance of the water in $Tank_1$.

### The Extended Stolen Car Problem

In the so-called "Extended Stolen Car Problem", Fred parks his car and then waits two intervals. We are told that at the end of that time, his car has been stolen. The representation of this problem in our formalism is straightforward. Note, however, that if one's axiomatization does not include the possibility of miracles, then the observation prunes *all* models (thus indicating that the axiomatization is too restrictive to allow any adequate description of the world to be built). If miracles are included in the axiomatization then two models will be produced—one in which the car is miraculously stolen during the first wait event, and one in which it is stolen during the second (with no preference between the two).

Now consider what happens if we enrich our representation so we can ask whether Fred is alive after the car is stolen. The provision of a fluent describing Fred's vitality seems a trivial enrichment of the problem. One would expect $Alive$ to simply persist (since there is no reason to believe it changes). Problems arise in Baker's formalism, however, precisely because there is no distinction between persistence abnormalities and modeling abnormalities. In that formalism, abnormalities are functions of situations. Thus the disappearance of the car in a situation in which Fred is alive is a distinct (and incomparable) abnormality from its disappearance in a situation in which he is dead. This results in an extra, unwanted, minimal model, in which Fred dies and then his car is stolen. This leaves Baker agnostic about Fred's health after the two *waits*.

Since modeling abnormalities are a function of the time at which situations occur, this enriched problem presents no difficulty for our formalism. The abnormality of having the car stolen during the second wait is the same regardless of the other fluent-values. Thus the model in which Fred dies and then has his car stolen has two modeling abnormalities, while the model in which he survives to suffer the indignity of losing his car has only one, and hence is preferred. We are thus left with the original pair of models (one in which the car is miraculously stolen during the first wait event, and one in which it is stolen during the second), and we predict unambiguously that Fred remains alive.

Now consider the case in which we are told that Fred

[9] The details, which are straightforward, are given in [Crawford & Etherington 1992].

[10] There are actually several other minimal models that are later eliminated and so do not materially affect the result.

$$\textbf{M1:}\ \overset{alive}{\neg stolen} \rightarrow \overset{\neg alive}{\neg stolen} \rightarrow \overset{\neg alive}{stolen}$$
$$\{Miracle(wait, t_1, alive), Miracle(wait, t_2, stolen)\}$$

$$\textbf{M2:}\ \overset{alive}{\neg stolen} \rightarrow \overset{alive}{\neg stolen} \rightarrow \overset{alive}{stolen}$$
$$\{Miracle(wait, t_2, stolen)\}$$

$$\textbf{M3:}\ \overset{alive}{\neg stolen} \rightarrow \overset{alive}{\neg stolen} \rightarrow \overset{\neg alive}{stolen}$$
$$\{Miracle(wait, t_2, stolen), Miracle(wait, t_2, alive)\}$$

$$\textbf{M4:}\ \overset{alive}{\neg stolen} \rightarrow \overset{alive}{stolen} \rightarrow \overset{alive}{stolen}$$
$$\{Miracle(wait, t_1, stolen)\}$$

$$\textbf{M5:}\ \overset{alive}{\neg stolen} \rightarrow \overset{alive}{stolen} \rightarrow \overset{\neg alive}{stolen}$$
$$\{Miracle(wait, t_1, stolen), Miracle(wait, t_2, alive)\}$$

$$\textbf{M6:}\ \overset{alive}{\neg stolen} \rightarrow \overset{\neg alive}{stolen} \rightarrow \overset{\neg alive}{stolen}$$
$$\{Miracle(wait, t_1, stolen), Miracle(wait, t_1, alive)\}$$

Figure 3: Six possible models of the extended stolen car problem (before minimization of $MAb$).

just cannot face the loss of his car—if the car is stolen during the first $Wait$ then Fred is dead after the second $Wait$. Since this statement concerns a particular time, our formalism treats it as an observation and uses it to prune possible qualitative models. Without loss of generality, we can restrict our attention to models with two or fewer $MAb$s and ignore differences between models that only affect states unreachable from the given initial state. This leaves eleven possible models. If we simulate each of these, and prune on the observation that the car is stolen after the two $Wait$s, then we are left with the six models shown in Figure 3 (the $MAb$s for each model are shown in braces). (Note that if we minimize $MAb$ at this point we get $M_2$ and $M_4$ as expected.)

We then prune using the observation that if the car is stolen during the first $Wait$ then Fred is dead after the second $Wait$. This eliminates only $M_4$. We can now minimize $MAb$, and eliminate $M_1$ and $M_3$ (whose $MAb$s are supersets of the $MAb$ in $M_2$). We are thus left with three possibilities—the car may be stolen during the second wait, in which case Fred survives, or it may be stolen during the first wait, in which case Fred dies during one of the waits. All of these are reasonable behaviors since the observation only stipulated that if the car is stolen then Fred is dead after the second wait—it did not require him to die *after* the car was stolen!

## Conclusions and Related Work

We have shown that preserving qualitative reasoning's distinction between model building and simulation avoids many of the problems faced by current nonmonotonic formalisms for reasoning about change. In particular, the problems (such as unwanted multiple models and unexplained persistence violations) typically caused by "unexpected" observations are avoided in our approach.

The motivation for our work is much like that for causal minimization [Lifschitz 1987; Lin & Shoham 1991]: no change should occur that is not explained by the model. Further, the separation of simulation from observations can be justified on causal grounds—observations about the system should not be allowed to cause violations of persistence (rather, they should be explained by an appropriate model). However, our approach appears to have no difficulties with ramifications (which cause problems for causal minimization).

In retrospect, we can see why other earlier approaches worked on some problems and failed on others. Baker's formalism [Baker 1989] was the first to adequately axiomatize simulation (by axiomatizing the existence of all situations and then "simulating" by circumscriptively determining the connections between situations through *result* links). However, Baker did not recognize model-building as a separate problem, and thus could not handle problems like the extended stolen car problem. Similarly, Morgenstern and Stein [1988] ignore the simulation problem by assuming a complete description of the effects of each action, but handle part of the model-building problem by showing how one can derive minimal sets of actions necessary to explain a given outcome.

The Yale Shooting Problem and its relatives have never been of much interest to the qualitative reasoning community. This is understandable since the model building/simulation distinction prevents them from occurring (in the normal model, waiting has no effect and the target dies, and in one possible abnormal model waiting unloads guns and he lives). However, the approach presented here provides a formal characterization of current QR practice, and hence a basis for formally studying extensions to that practice that address several outstanding QR problems.

First, dealing with *region transitions*—points at which a qualitative model ceases to describe a system and a new model must be built (*e.g.*, because a pipe breaks, or a tank overflows)—has recently emerged as a tricky problem in QR [Forbus 1989; Sandewall, 1989; Rayner, 1991]. The semantics of region transitions is poorly understood (especially since, in commonsense reasoning, region transitions often involve instantaneous changes in the value of otherwise smooth functions), and may well require a formal theory of change.

Second, the paradigm of building the entire state transition graph (showing each possible state and the effect of each possible action thereon) roughly corresponds to the QR notion of a total envisionment [Forbus 1984]. In the interest of computational efficiency, one might want to modify the formalism so that the initial fluent values, and statements about which actions occur when, *are* given to the simulator. This would allow the simulator to build only a part of the

state transition graph. Such an approach would correspond to the QR idea of an incremental envisionment [Crawford, Farquhar & Kuipers 1990].

It is generally believed that the only differences between total and incremental envisionments are computational, but the exact conditions under which the approaches are semantically equivalent have not been studied. Further, showing that algorithms based on incremental envisionments produce the same simulations as algorithms based on total envisionments has proved difficult [Forbus 1992]. This is especially true in domains, such as qualitative physics, in which the existence of some conceptual objects may be conditioned on modeling assumptions (*i.e.*, objects may exist in some qualitative models but not in others). Our formalism could provide a formal basis for studying such equivalences.

Finally, there is a debate in the qualitative reasoning community concerning whether approaches that separate model building and simulation need to be augmented with explicit representations of causality. This work indicates that one can get surprisingly far in a formalism without causality, but the role of and necessity for causal information remain open questions. One interesting avenue for further work is to study whether there is a fundamental difference between formalisms based on the modeling/simulation distinction and formalisms that explicitly axiomatize causality. Our experience with this formalism suggests that the separation of model building and simulation may capture many aspects of causation.

## Acknowledgments

We would like to thank Ben Kuipers for several fruitful discussions on this topic, and for comments on an earlier draft of this paper.

## References

Baker, A.B. (1989). A simple solution to the Yale Shooting Problem. *Proc. First Int'l Conf. on Principles of Knowledge Representation and Reasoning.* pp. 11–20.

Crawford, J.M. and Etherington, D.W. (1992). *Formalizing Reasoning About Change: A Qualitative Reasoning Approach.* In preparation.

Crawford, J., Farquhar, A., and Kuipers, B. (1990). QPC: A compiler from physical models into Qualitative Differential Equations. *Proc. Eighth Nat'l Conf. on Artificial Intelligence.* pp 365–372.

D. Dvorak and B. J. Kuipers. (1989). Model-based monitoring of dynamic systems. *Proc. Eleventh Int'l Joint Conf. on Artificial Intelligence.* pp. 1238–1243.

Forbus, K.D. (1984). Qualitative process theory. *Artificial Intelligence*, 24:85–168, 1984.

Forbus, K.D. (1989). Introducing actions into qualitative simulation. *Proc. Eleventh Int'l Joint Conf. on Artificial Intelligence.* pp. 1273–1278.

Forbus, K.D. (1992). Personal communication.

Ginsberg, M.L. (ed.) (1987). *Readings in Nonmonotonic Reasoning.* Morgan Kaufmann, Los Altos, CA.

Kautz, H.A. (1986). The logic of persistence. *Proc. Fifth Nat'l Conf. on Artificial Intelligence.* pp. 401–405.

Kuipers, B.J. (1986). Qualitative simulation. *Artificial Intelligence*, 29:289–338.

Lifschitz, V. (1987). Formal theories of action. In Matthew L. Ginsberg, ed., *Readings in Nonmonotonic Reasoning.* Morgan Kaufmann, Los Altos, CA. pp. 410–432.

Lifschitz, V., and Rabinov, A. (1989). Things that change by themselves. In *Proc. Eleventh Int'l Joint Conf. on Artificial Intelligence.* pp. 864–867.

Lin, F. and Shoham Y. (1991). Provably correct theories of action (preliminary report). *Proc. 9th Nat'l Conf. on Artificial Intelligence.* pp. 349–354.

McCarthy, J., and Hayes, P.J. (1969) Some Philosophical Problems from the Standpoint of Artificial Intelligence. In B. Meltzer and D. Mitchie (eds.) *Machine Intelligence 4.* Edinburgh University Press, pp. 462–502.

Morgenstern, L., and Stein, L.A. (1988) Why things go wrong: A formal theory of causal reasoning. *Proc. Seventh Nat'l Conf. on Artificial Intelligence.* pp. 518–523.

Rayner, M. (1991) On the applicability of nonmonotonic logic to formal reasoning in continuous time. *Artificial Intelligence*, 49:345–360.

Sandewall, E. (1989) Combining logic and differential equations for describing real-world systems. *Proc. First Int'l Conf. on Principles of Knowledge Representation and Reasoning.* pp. 412–420.

Weld, D.S., and de Kleer, J. (1990). *Readings in Qualitative Reasoning About Physical Systems.* Morgan Kaufmann, Los Altos, CA, 1990.