Belief Revision with Unreliable Observations

Craig Boutilier[*]
Dept. Computer Science
University of British Columbia
Vancouver, British Columbia
Canada, V6T 1W5
cebly@cs.ubc.ca

Nir Friedman[†‡]
Computer Science Division
387 Soda Hall
University of California
Berkeley, CA 94720
nir@cs.berkeley.edu

Joseph Y. Halpern[§]
Dept. Computer Science
Cornell University
Ithaca, NY 14850
halpern@cs.cornell.edu

Abstract

Research in belief revision has been dominated by work that lies firmly within the classic AGM paradigm, characterized by a well-known set of postulates governing the behavior of "rational" revision functions. A postulate that is rarely criticized is the *success postulate*: the result of revising by an observed proposition φ results in belief in φ. This postulate, however, is often undesirable in settings where an agent's observations may be imprecise or noisy. We propose a semantics that captures a new ontology for studying revision functions, which can handle noisy observations in a natural way while retaining the classical AGM model as a special case. We present a characterization theorem for our semantics, and describe a number of natural special cases that allow ease of specification and reasoning with revision functions. In particular, by making the *Markov assumption*, we can easily specify and reason about revision.

1 Introduction

The process by which an agent revises its beliefs when it obtains new information about the world, that is, the process of *belief change*, has been the focus of considerable study in philosophy and artificial intelligence. One of the best known and most studied theories of belief change is the classic *AGM theory of belief revision* of Alchourrón, Gärdenfors and Makinson [2, 18]. Recent years have seen many extensions and refinements of the AGM paradigm, including the distinction between *revision* and *update* [23, 32], the proposal of models that combine the two [8, 15], and the acceptance of the notion that epistemic states are much richer than simple belief sets [4, 17, 27, 29].

All of these advances can be viewed as refinements of the AGM paradigm, for none contradict the basic, if—in retrospect—somewhat limited, view of revision profferred by AGM. However, as noted in [14], there are reasons to question some of their rationality postulates, even some that have been viewed as "beyond controversy". One example is the *success postulate*, which asserts that when an agent revises its beliefs based on new information φ, the resulting belief set should contain φ; that is, the revision should "succeed" in the incorporation of the new information. Generally, this requires that, in order to accept φ, the agent give up some of its old beliefs to remain consistent.

As argued in [14], to justify the success postulate (or any other postulate), we must carefully consider the particular process we hope to characterize as well as the ontology adopted in that characterization. Gärdenfors [18] provides one interpretation of belief revision for which the success postulate is appropriate. Under this interpretation, the agent's beliefs consist of those propositions the agent accepts as being true and the agent revises by φ only if it accepts φ as being true. In this case, the success postulate holds almost by definition.

In much work on revision, it is implicitly assumed that an agent should revise by φ after observing φ. The reasonableness of this assumption depends in part on the language being used. For example, if the agent is a robot making observations and φ talks about the reading of a sensor—for example, saying that a particular sensor had a high reading—then the success postulate may again be deemed acceptable. Of course, the relationship between the propositions that talk about the robot's sensors and more interesting propositions that talk about what is actually true in the "external" world must be modeled if the robot is to draw useful inferences [8]. The relationship will generally be complicated because of sensor noise, unreliability, and so on. One may instead wish to model a situation of this type by assuming the robot can directly observe the truth values of external propositions, but that these direct observations may be corrupted. Adopting this ontology, the success postulate is no longer reasonable: an observed proposition may contradict such strongly held beliefs that the robot has no choice but to dismiss the observation as incorrect.[1]

As another example where the success postulate may be questionable, imagine an agent conducting a market survey

[*]This work was supported by NSERC Research Grant OGP0121843 and IRIS-II Project IC-7, and was undertaken while the author was visiting the University of Toronto. Thanks also to the generous support of the Killam Foundation.

[†]Current address: Institute of Computer Science, The Hebrew University, Givat Ram, Jerusalem 91904, Israel. nir@cs.huji.ac.il

[‡]Work was supported in part by ARO under the MURI program "Integrated Approach to Intelligent Systems", grant number DAAH04-96-1-0341.

[§]Work was supported in part by NSF under grant IRI-96-25901 and by the Air Force Office of Scientific Research under grant F49620-96-1-0323.

[1]As we shall see, "dismiss" is too strong a word, for an observation that is not incorporated into the agent's belief set will still have an impact on its epistemic state, for instance, by predisposing it to the future acceptance of that proposition.

by having people fill in on-line questionnaires. By sending several different questionnaires to the same person, the agent can obtain multiple observations of, say, the person's salary. None of these observations may be the person's actual salary; in fact, the agent might believe that most people tend to exaggerate their salaries when filling out such questionnaires.

Notice also that the success postulate imposes an overwhelming bias toward accepting the most recently observed propositions. An agent that observes a sequence of propositions containing some number of φs and $\neg\varphi$s is bound to accept φ or $\neg\varphi$ depending on which it observed most recently, regardless of any feature of the actual history of observations. For instance, the robot will ignore the fact that twice as many φ observations as $\neg\varphi$ observations were made; and our survey agent must ignore the fact that men tend to inflate their reported salaries over time and accept the most *inaccurate* observation. Clearly what is required is a model of revision that lets us understand when the success postulate is reasonable and allows us to discard it when it is not.

In this paper, we propose a model of belief revision that deals with imprecise observations by adopting the second ontology mentioned above. We assume that an agent has access to a stream of observed propositions, but that it is under no obligation to incorporate any particular observed proposition into its belief set. Generally, a proposition will be accepted only if the likelihood that the proposition is true given the agent's current sequence of observations "outweighs" the agent's prior belief that it was false. The basic intuitions are drawn from the standard Bayesian model of belief change. Roughly speaking, once we have represented in our model the correlation between a given stream of observations and the truth of various propositions, we can simply condition on the observations. The key point is that conditioning on the event of *observing* φ is very different from conditioning on the event φ. After conditioning on φ, the probability of φ is 1; after conditioning on observing φ, the probability of φ depends on the prior probability of φ and the correlation between the observation and φ actually being true.

To use these ideas in the context of belief revision, we must use a more qualitative measure of uncertainty than probability—here we adopt Spohn's [31] *ranking functions*. Nevertheless, the basic intuitions are drawn from the standard probabilistic approach. Indeed, it is remarkable how little work is needed to apply these intuitions in a qualitative setting. This emphasizes how small the gap is between belief revision and probability kinematics. We note, however, that our model differs from qualitative adaptations of Jeffrey's Rule [22] devised for belief revision [11, 19, 31] (see Section 2 for further discussion).

The rest of this paper is organized as follows. In Section 2, we discuss the AGM model, and describe *generalized revision functions* for dealing with sequences of observations. In Section 3, we present our basic framework, which is taken from Friedman and Halpern [16, 17]. We define *observation systems* that allow unreliable observations, show how conditioning can be used to effect belief revision, and characterize the class of generalized revision functions determined by observation systems. In Section 4, we consider the important class of observation systems that satisfy the *Markov assumption*, allowing revision functions to be specified concisely and naturally. In Section 5, we consider two further special cases where (a) observations are more likely to be true than false; and (b) observations are known to be accurate. We conclude with a brief discussion of related and future work.

2 The AGM Theory of Belief Revision

Throughout, we assume that an agent has a deductively closed *belief set* K, a set of sentences drawn from some logical language reflecting the agent's beliefs about the current state of the world. For ease of presentation, we assume a classical propositional language, denoted $\mathcal{L}$, and consequence operation *Cn*. The belief set K will often be generated by some finite knowledge base *KB* (i.e., $K = Cn(KB)$). The identically true and false propositions are denoted $\top$ and $\bot$, respectively. Given a set of possible worlds W and $\varphi \in \mathcal{L}$, we denote by $[\![\varphi]\!]$ the set of φ-worlds, the elements of W satisfying φ.[2]

Given a belief set K, an agent will often obtain information φ not present in K. In this case, K must be *revised* to incorporate φ. If φ is consistent with K, one expects φ to simply be added to K. More problematic is the case when $K \vdash \neg\varphi$; certain beliefs must be given up before φ is adopted. The *AGM theory* provides a set of postulates governing this process. We use K^*_φ to denote the *revision* of K by φ. Of interest here is the following:

(R2) $\varphi \in K^*_\varphi$.

R2 is the success postulate mentioned in the introduction; it says that φ is believed after revising by φ. We refer the reader to [18] for the remaining postulates and a discussion of the AGM theory.

Unfortunately, while the postulates constrain possible revisions, they do not dictate the precise beliefs that should be retracted when φ is observed. An alternative model of revision, based on the notion of *epistemic entrenchment* [18], has a more constructive nature. Given a belief set K, we can characterize the revision of K by ordering beliefs according to our willingness to give them up. If one of two beliefs must be retracted in order to accommodate some new fact, the less entrenched belief will be relinquished, while the more entrenched persists.

Semantically, an entrenchment relation (hence a revision function) can be modeled by associating with each set of possible worlds a plausibility, in any of a number of ways [5, 10, 17, 20]. For the purposes of this paper, we adopt Spohn's *ordinal conditional functions* or κ-rankings [19, 31]. A function $\kappa : W \to \mathbf{N} \cup \{\infty\}$ assigns to each world a ranking reflecting its plausibility: if $\kappa(w) < \kappa(v)$ then w is more plausible than v. We insist that $\kappa^{-1}(0) \neq \emptyset$, so that maximally plausible worlds are assigned rank 0. If $\kappa(w) = \infty$, we say w is *impossible*. If $U \subseteq W$, then $\kappa(U) = \min_{u \in U} \kappa(u)$.

Following [4, 16, 27, 29], we distinguish the agent's *epistemic state* from its belief set. We define the form of the epistemic state carefully in Section 3. For now we simply require that it includes a ranking κ. This ranking then determines the agent's belief set K as follows:

$$K = \{\varphi \in \mathcal{L} : \kappa^{-1}(0) \subseteq [\![\varphi]\!]\}. \qquad (1)$$

Thus, the formulas in K are precisely those that are true in all worlds of rank 0.

The ranking κ also induces a revision function: to revise by φ an agent adopts the most plausible φ-worlds as epistemically possible. Thus, using $\min(\varphi, \kappa)$ to denote this set, we have

$$K^*_\varphi = \{\psi \in \mathcal{L} : \min(\varphi, \kappa) \subseteq [\![\psi]\!]\}$$

[2] In our setting, we can safely identify the possible worlds with valuations over $\mathcal{L}$, although in general we must distinguish the two.

If $[\![\varphi]\!] \cap W = \emptyset$, we set $\min(\varphi, \kappa) = \emptyset$ and $K^*_\varphi = \mathcal{L}$ (the inconsistent belief set). It is normally assumed that $[\![\varphi]\!] \cap W \neq \emptyset$ for every satisfiable φ — thus every satisfiable proposition is accorded some degree of plausibility. It is well-known that this type of model induces the class of revision functions sanctioned by the AGM postulates [5, 19, 20].

We define *conditional plausibility*, for $U, V \subseteq W$ and $\kappa(U) \neq \infty$, as:

$$\kappa(V|U) = \kappa(V \wedge U) - \kappa(U).$$

Intuitively, this denotes the degree to which V would be considered plausible if U were believed.

These notions are strongly reminiscent of standard concepts from probability theory. Indeed, the role of $+$ in probability is assumed by min in the theory of rankings, while the role of $\times$ is assumed by $+$ (so, in the definition of conditioning, division becomes subtraction). In fact, a κ-ranking can be interpreted as a semi-qualitative probability distribution. Using the ε*-semantics* of Adams [1], Goldszmidt and Pearl [19] show how one can interpret the κ values of propositions as "order of magnitude" probabilities.

It has been remarked by a number of authors that models of revision based on epistemic entrenchment or κ-rankings are not strong enough to adequately capture *iterated revision* [4, 16, 27, 29]. Specifically, while these models determine the content of a new belief set when φ is observed, given an epistemic state, they do not determine the new epistemic state (or ranking) associated with the new belief set. To deal with iteration semantically, we need a way of determining a new epistemic state, given an observation [7, 9, 26, 30]. Spohn's *conditioning* operation [31] does just this. When an observation φ is made, all $\neg\varphi$-worlds are deemed impossible and removed from the ranking (or set to ∞). The remaining φ-worlds retain their relative plausibilities, with the resulting ranking κ^*_φ renormalized; formally we have

$$\kappa^*_\varphi(w) = \begin{cases} \kappa(w) - \kappa(\varphi) & \text{if } w \models \varphi \\ \infty & \text{if } w \not\models \varphi. \end{cases} \qquad (2)$$

Thus each observation determines not just a revised belief set, but a new epistemic state which can be used to model subsequent revisions.

Spohn also proposed a more general model of revision called α*-conditioning*. Rather than accepting an observed proposition φ with certainty, φ is accepted with degree α, with $\neg\varphi$-worlds retaining a certain plausibility. This model can be viewed as a way of dealing with noisy observations (and has been developed further in [11, 19]). In fact, this model is a qualitative analogue of *Jeffrey's Rule* [22] for probabilistic belief update. Jeffrey's rule is a generalization of conditioning where a piece of evidence can be accepted with a given probability.

Goldszmidt and Pearl [19] argue that Jeffrey's rule is unreasonable since it requires that the observation φ is associated with the agent's posterior degree of belief in φ. As an alternative, they propose *L-conditioning*, where the strength associated with observing φ conveys the difference in the *evidential support* that the observation gives to worlds that satisfy φ and to worlds that satisfy $\neg\varphi$. They use a qualitative version of Bayes' rule that combines this support with the agent's prior ranking of φ and $\neg\varphi$ to get a *posterior* ranking over both propositions. Then they apply Jeffrey's rule to update the ranking over worlds to match this posterior.

The approach we propose is different from and, in a sense, more general, than both of these qualitative update rules. Unlike Jeffrey's rule, we do not assume that there is any doubt that φ has been observed but, as we said earlier, we distinguish observing φ from φ being true. Like Goldszmidt and Pearl's approach, our approach relies on Bayes rule to combine the evidence with a prior rankings. However, unlike their approach, we assume that the evidence provides support for each possible world (rather than to the propositions φ and $\neg\varphi$), and thus we do not have to appeal to Jeffrey's rule. In this sense, our proposal has more in common with probabilistic *observation models* that are standard in decision and control theory [3, 24].

A general way of thinking about iterated revision is not to think of revision functions as mapping from (belief set, observation) pairs to belief sets, but as mapping from finite observation sequences to belief sets. More precisely, assume that an agent's observations are drawn from language $\mathcal{L}$. We use $\langle \varphi_1, \varphi_2, \cdots, \varphi_n \rangle$ to denote the length n sequence consisting of $\varphi_1, \varphi_2, \ldots$; and $\langle\rangle$ denotes the length-0 sequence. Let $\mathcal{O}$ denote the set of all finite sequences of observations, and let $\mathcal{B}$ denote the set of all belief sets over $\mathcal{L}$.

Definition 2.1: A *generalized revision function* B is a mapping $B : \mathcal{O} \to \mathcal{B}$.

This definition deals naturally with iterated revision. Furthermore, there is no need to specify an initial belief set: the agent's prior beliefs are captured by the belief set $B(\langle\rangle)$.[3]

3 An Ontology for Imprecise Observations

3.1 Observation Systems

In this section, we present a framework that allows us to describe what is true in the world, what the agent observes, and the plausibility of these observations. The framework is essentially that of Friedman and Halpern [17, 16], which in turn is based on the multi-agent framework of [21]. We briefly review the details here; further discussion and motivation can be found in [12].

The key assumption in the multi-agent system framework is that we can characterize the system by describing it in terms of a *state* that changes over time. Formally, we assume that at each point in time, the agent is in some *local state*. Intuitively, this local state encodes the information the agent has observed thus far and its ranking. There is also an *environment*, whose state encodes relevant aspects of the system that are not part of the agent's local state. In this case, the relevant information is simply the *state of the world*. A *global state* (s_e, s_a) consists of the environment state s_e and the local state s_a of the agent. A *run* of the system is a function from time (which here ranges over $\mathbf{N}$) to global states. Thus, if r is a run, then $r(0), r(1), \ldots$ is a sequence of global states that completely describes a possible system execution. A *system* consists of a set of runs that dictates all the possible behaviors of the system.

Given a system $\mathcal{R}$, we refer to a pair (r, m) consisting of a run $r \in \mathcal{R}$ and a time m as a *point*. If $r(m) = (s_e, s_a)$, we define $r_a(m) = s_a$ and $r_e(m) = s_e$. A *(ranked) interpreted system* is a tuple $\mathcal{I} = (\mathcal{R}, \kappa, \pi)$, consisting of a system $\mathcal{R}$, a ranking κ on the runs in r, and an *interpretation* π, which associates with each point a truth assignment for $\mathcal{L}$. The ranking κ represents the agent's initial ranking of runs. Notice that in the previous section, the agent simply ranks possible worlds; here the agent ranks the relative plausibility

[3] Our definition of generalized revision functions is similar to that of Lehmann [26].

of entire "evolutions" of both its local state and the environment state. Of course, in general, it is infeasible for the agent to come up with a complete ranking over all possible runs. Later, we discuss some simplifying assumptions that make obtaining such a ranking more feasible.

To capture belief revision, we consider a special class of interpreted systems called *observation systems*.[4] We assume that the agent makes observations, which are characterized by formulas in $\mathcal{L}$, and that its local state consists of the sequence of observations that it has made. We also assume that the environment state is a truth assignment for $\mathcal{L}$, reflecting the actual state of the world. As observed by Katsuno and Mendelzon [23], the AGM postulates assume that the world is *static*; to capture this, we require that the environment state does not change over time. An *observation system* (OS) is a ranked interpreted system $(\mathcal{R}, \kappa, \pi)$ that satisfies the following three assumptions for every point (r, m):

- The environment state $r_e(m)$ is a truth assignment to the formulas in $\mathcal{L}$ that agrees with π at (r, m) (that is, $\pi(r, m) = r_e(m)$), and $r_e(m) = r_e(0)$.
- The agent's state $r_a(m)$ is a sequence of the form $\langle \varphi_1, \ldots, \varphi_m \rangle$ of formulas in $\mathcal{L}$; and if $m \geq 1$ and $r_a(m) = \langle \varphi_1, \ldots, \varphi_m \rangle$, then $r_a(m-1) = \langle \varphi_1, \ldots, \varphi_{m-1} \rangle$.
- If $\vec{\varphi}$ is a sequence $\langle \varphi_1, \ldots, \varphi_m \rangle$ of observations such that $r_a(m) = \vec{\varphi}$ for some run r, then $\kappa([\![Obs(\vec{\varphi})]\!]) \neq \infty$, where $[\![Obs(\vec{\varphi})]\!] = \{r' : r'_a(m) = \vec{\varphi}\}$. Intuitively, this says that any sequence of observations that actually arises in the system is initially considered possible.

Notice that the form of the agent's local state makes explicit an important implicit assumption: that the agent remembers all of its previous observations. Its local state "grows" at every step.[5]

We introduce the following notation before proceeding. Since the environment state is fixed throughout a given run r, we use r_e to denote this state, dropping the time index from $r_e(m)$. We use the notation $[\![C]\!]$ to denote the set of runs in $\mathcal{R}$ that satisfy condition C. In particular:

- For any length-m observation sequence $\vec{\varphi}$, $[\![Obs(\vec{\varphi})]\!] = \{r : r_a(m) = \vec{\varphi}\}$ denotes the set of runs in $\mathcal{R}$ in which $\vec{\varphi}$ is the initial sequence of observations.
- For any $\psi \in \mathcal{L}$, $[\![Obs^m(\psi)]\!] = \{r : r_a(m) = \vec{\varphi} \cdot \psi,$ for some length-$(m-1)$ sequence $\vec{\varphi}\}$ is the set of runs in which ψ is the mth observation.
- For any $\psi \in \mathcal{L}$, $[\![\psi]\!] = \{r : r_e \models \psi\}$ is the set of runs in which the (fixed) environment state satisfies ψ
- For any truth assignment w, $[\![w]\!] = \{r : r_e = w\}$ is the set of runs whose (fixed) environment state is w.
- For any length-m sequence $\vec{\varphi}$ and truth assignment w, $[\![w, Obs(\vec{\varphi})]\!] = \{r : r_e = w, r_a(m) = \vec{\varphi}\}$.

We stress again the difference between $[\![Obs^m(\psi)]\!]$ and $[\![\psi]\!]$. The former is the event of observing ψ at time m; the latter is the event of ψ being true.

In the analysis of the AGM framework in [16], an extra requirement is placed on OSs: it is required that, in any run r, if $r_a(m) = \langle \varphi_1, \ldots, \varphi_m \rangle$, then $\varphi_1 \wedge \ldots \wedge \varphi_m$ is true according to the truth assignment at $r_e(m)$. This requirement forces the observations to be accurate; if φ is observed, then it must be true of the world. It is precisely this requirement that we drop here to allow for noisy observations. The initial κ ranking specifies (among other things) the likelihood of such noisy observations.

We can now associate with each point (r, m) a ranking $\kappa^{r,m}$ on the runs. We take $\kappa^{r,0} = \kappa$, and define

$$\kappa^{r,m+1}(U) = \kappa^{r,m}(U \mid [\![Obs(r_a(m+1))]\!])$$

for each subset U of runs. Thus, $\kappa^{r,m+1}$ is the result of conditioning $\kappa^{r,m}$ on the observations the agent has made up to the point $(r, m+1)$. Because the agent has perfect recall, it is easy to see that conditioning on the sequence of observations $r_a(m+1)$ is equivalent to conditioning on the last observation φ'. More precisely,

Lemma 3.1: *If $r_a(m+1) = \langle \varphi_1, \cdots, \varphi_{m+1} \rangle$, then*

$$\kappa^{r,m+1}(U) = \kappa^{r,m}(U \mid [\![Obs^{m+1}(\varphi)]\!]).$$

It is immediate from the definition that $\kappa^{r,m}$ depends only on the agent's local state $r_a(m)$; if $r_a(m) = r'_a(m)$, then $\kappa^{r,m} = \kappa^{r',m}$. Thus, we usually write $\kappa^{\vec{\varphi}}$ to denote the ranking $\kappa^{r,m}$ such that $r_a(m) = \vec{\varphi}$. We take the agent's *epistemic state at the point* (r, m) to consist of its local state $r_a(m)$ and the ranking $\kappa^{r,m}$. Since the ranking is determined by the local state, we can safely identify the agent's epistemic state with its local state. We note that we can generalize our model by embedding the agent's κ-rankings in the local state without difficulty, allowing different initial rankings in different situations. For simplicity of exposition, we consider only a fixed ranking.

The beliefs an agent holds about the world at any point (r, m) are determined by the runs it considers most plausible at that point. Mirroring (1), we define the agent's belief set $Bel(\mathcal{I}, r, m)$ at point (r, m) in a system $\mathcal{I}$ as

$$Bel(\mathcal{I}, r, m) = \{\varphi \in \mathcal{L} : ((\kappa^{r,m})^{-1}(0) \subseteq [\![\varphi]\!]\}. \quad (3)$$

Again, notice that an agent's belief set depends only on its local state; that is, if $r_a(m) = r'_a(m')$, then $Bel(\mathcal{I}, r, m) = Bel(\mathcal{I}, r', m')$. However, it may well be that the agent has the same belief set at two points where it has quite distinct local states; moreover, revisions of these belief sets can proceed differently. Thus, an observation system $\mathcal{I}$ defines a generalized revision function that maps epistemic states to belief sets:

$$B_{\mathcal{I}}(\vec{\varphi}) = \begin{cases} Bel(\mathcal{I}, r, m) & \text{for } (r, m) \text{ such that } r_a(m) = \vec{\varphi} \\ Cn(\bot) & \text{if } \kappa([\![Obs(\vec{\varphi})]\!]) = \infty. \end{cases}$$

Example 3.2: As an example of an OS, consider the marketing survey example discussed in the introduction. Suppose our marketing agent sends three different surveys to one person. In each of them, the respondent must mark his salary, in multiples of ten thousand. Initially, the agent considers the person's annual salary to be either \$30,000, \$40,000, or \$50,000, each equally plausible. The agent also knows how plausible various observation sequences for these three surveys are—if the person's salary is \$10,000x, he will report one of the following sequences:

- $\langle x+1, x+2, x+3 \rangle$: the incremental exaggerator
- $\langle x+2, x+2, x+2 \rangle$: the consistent exaggerator

[4] Observation systems are a special case of the *belief change systems* considered in [13, 16].

[5] This is analogous to the assumption of *perfect recall* in game theory [28] and distributed computing [12].

- $\langle x, x+1, x+1 \rangle$: the reluctant exaggerator
- $\langle x, x, x \rangle$: the non-exaggerator.

The agent considers it most likely that the survey recipient is an incremental or consistent exaggerator, less likely that he is a reluctant exaggerator, and quite implausible that he is a non-exaggerator. The ranking of any run with environment state (salary) $x \in \{3, 4, 5\}$ and sequence of survey answers $\langle x_1, x_2, x_3 \rangle$ is: 0 if the sequence follows the incremental or consistent pattern (given state x); 1 if it follows the reluctant pattern; 2 if it is unexaggerated; and 3 if it follows any other pattern. In addition, for $x \notin \{3, 4, 5\}$, we set κ of any run r with $r_e = x$ to be 3.

In the resulting system $\mathcal{I}$, the agent's initial belief set, $B_{\mathcal{I}}(\langle\rangle)$, is $Cn(x \in \{3, 4, 5\})$. The agent's belief set after getting a response of 6 to the first survey, $B_{\mathcal{I}}(\langle 6 \rangle)$, is $Cn(x \in \{4, 5\})$. This observation rules out the most plausible runs where the agent's actual salary is \$30,000, since they are incompatible with the agent being an incremental or consistent exaggerator. After then observing 7, the agent's belief set is $B_{\mathcal{I}}(\langle 6, 7 \rangle) = Cn(x = 5)$; he believes that he is dealing with an incremental exaggerator. Finally, if he then observes 7 again, his belief set is $B_{\mathcal{I}}(\langle 6, 7, 7 \rangle) = Cn(x = 6)$; he believes that he is dealing with a reluctant exaggerator. ∎

3.2 Expressive Power

We now examine properties of the revision functions induced by OSs, and the expressive power of OSs. We might ask whether the (ordinary) revision function determined by $B_{\mathcal{I}}$ satisfies the AGM postulates. The answer is, of course, no. Not surprisingly, the success postulate is not satisfied, for an agent may observe ψ and still believe $\neg\psi$ (as illustrated in our example above).

With respect to expressive power, we might ask whether all possible generalized revision functions (mapping observation sequences to belief sets) can be represented by OSs. This is not the case in general, but a particular class of revision functions does correspond to OSs.

We say that an observation is *unsurprising* if it does not cause the agent to retract any of its previous beliefs. That is, its belief set after the observation is a superset of its prior belief set. We impose the following rationality postulate on generalized revision functions:

(O1) For any finite observation sequence $\vec{\varphi}$ there exists a nonempty set of observations $Plaus(\vec{\varphi})$ such that

$$Cn(\cap\{B(\vec{\varphi} \cdot \psi) : \psi \in Plaus(\vec{\varphi})\}) = B(\vec{\varphi}).$$

According to O1, for every observation sequence $\vec{\varphi}$, there is a set $Plaus(\vec{\varphi})$ of observations, each of which is unsurprising with respect to the belief set $B(\vec{\varphi})$. To see this, note that O1 implies that if $\psi \in Plaus(\vec{\varphi})$, then $B(\vec{\varphi}) \subseteq B(\vec{\varphi} \cdot \psi)$, that is, the agent retains all of its beliefs after observing ψ. Moreover, this set of unsurprising observations "covers" the possibilities embodied by the belief set, in that any formula consistent with $B(\vec{\varphi})$ must be consistent with $B(\vec{\varphi} \cdot \psi)$ for some $\psi \in Plaus(\vec{\varphi})$.

With O1, we now can show the desired characterization theorems.

Theorem 3.3: *For any OS $\mathcal{I}$, the revision function $B_{\mathcal{I}}$ induced by $\mathcal{I}$ satisfies O1.*

Theorem 3.4: *Let B be an revision function satisfying O1. There exists an OS $\mathcal{I}$ such that $B = B_{\mathcal{I}}$.*

This shows that κ-rankings over runs are sufficient to represent any coherent revision function (i.e., satisfying O1) and thus can be viewed as a suitable semantics for revision based on unreliable observations.

4 Markovian Observation Models

To this point, we have placed few restrictions on the initial κ ranking on runs. This generality can cause several problems. First, the specification of such a κ ranking can be onerous, requiring (potentially) the individual ranking of all possible observation histories with respect to all possible truth assignments. Second, maintaining and updating such an explicit model imposes severe computational demands. Fortunately, there are a number of natural assumptions that can be made about the form of the observation model that make both the specification and reasoning tasks much more tractable.

Very often the state of the world completely determines the plausibility of various observations at a given point in time. In such a case, the history of past observations is irrelevant to the determination of the plausibility of the next observation if the state of the world is known. For instance, our agent conducting market surveys may know that a respondent provides independent salary reports at different points in time, the plausibility of a particular report being determined solely by the respondent's actual salary, not by their previous reports. In such a case, the "exaggeration patterns" described above no longer make sense. Instead the agent might assess the plausibility of a respondent reporting $x + k$ given that his salary is x.

We say an OS $\mathcal{I} = (\mathcal{R}, \kappa, \pi)$ is *Markovian* if it captures this intuition, which can be expressed formally as follows.

Definition 4.1: The OS $\mathcal{I} = (\mathcal{R}, \kappa, \pi)$ is *Markovian* if

(a) the likelihood of observing φ is independent of history and depends only on the state of the world, i.e., for all m, length-m sequences $\vec{\psi}$, and worlds w, we have

$$\kappa([\![Obs^{m+1}(\varphi)]\!] \mid [\![w, Obs(\vec{\psi})]\!]) = \kappa([\![Obs^{m+1}(\varphi)]\!] \mid [\![w]\!]).$$

(b) the likelihood of observing φ is independent of time, given the state of the world, i.e., for all m and m', we have

$$\kappa([\![Obs^{m}(\varphi)]\!] \mid [\![w]\!]) = \kappa([\![Obs^{m'}(\varphi)]\!] \mid [\![w]\!]).$$

The Markov assumption is standard in the probabilistic literature and has been argued to be widely applicable [25]. It is also adopted implicitly in much work in reasoning about action, planning, control and probabilistic inference with respect to system dynamics. Although it plays a key role in the observation models adopted in control theory and probabilistic reasoning [3, 24], it has received little attention in this respect within the qualitative planning literature.

The Markov assumption is also very powerful, allowing us to specify a ranking over runs relatively compactly. We need specify only two components: a *prior ranking* over worlds, i.e., $\kappa([\![w]\!])$ for each truth assignment w, and a family of *conditional observation rankings* of the form $\kappa([\![Obs^{\star}(\varphi)]\!] | [\![w]\!])$ for any observation φ and world w (we use the $\star$ to indicate that the observation plausibility is independent of m). Note that our conditional observation rankings differ dramatically from the general model, requiring that we rank individual observations, not infinite sequences of observations. These two components, however, determine the κ-ranking over runs.

Lemma 4.2: *Let $\mathcal{I} = (\mathcal{R}, \kappa, \pi)$ be a Markovian OS. Then the plausibility of the run $(w, \langle \varphi_1, \varphi_2, \ldots \rangle)$ is given by*

$$\kappa(w, \langle \varphi_1, \varphi_2, \cdots \rangle) = \kappa([\![w]\!]) + \sum_{j=1}^{\infty} \kappa([\![Obs^\star(\varphi_j)]\!] | [\![w]\!]). \tag{4}$$

Note that the infinite sum in the lemma may be ∞; this simply means that the run $(w, \langle \varphi_1, \varphi_2, \cdots \rangle)$ is impossible.

We can also easily characterize the ranking of an agent who has observed $\vec{\psi} = \langle \psi_1, \ldots, \psi_m \rangle$ at time m:

$$\begin{aligned} &\kappa^{\vec{\psi}}(w, \langle \varphi_1, \varphi_2, \cdots \rangle) \\ &= \begin{cases} \infty \text{ if } \vec{\psi} \neq \langle \varphi_1, \ldots, \varphi_m \rangle \\ \kappa^{\vec{\psi}}([\![w]\!]) + \sum_{j=m+1}^{\infty} \kappa([\![Obs^\star(\varphi_j)]\!] | [\![w]\!]) \text{ otherwise.} \end{cases} \end{aligned} \tag{5}$$

Thus, after observing $\vec{\psi}$, the agent's posterior over runs retains its Markovian structure, except that instead of using $\kappa([\![\alpha]\!])$, the prior over truth assignments, the agent now uses its posterior over truth assignments.

Example 4.3: We now reexamine the survey example. It is easy to verify that the system described in the previous example is not Markovian.[6] Instead, imagine that the agent believes a respondent with salary $\$x$ is most likely reply $x + 2$, then $x + 1$ then x, regardless of the number of times they are questioned. This can be modeled in a Markovian OS by (e.g.) assessing $\kappa([\![Obs^\star(x+2)]\!] | [\![x]\!]) = 0$, $\kappa([\![Obs^\star(x+1)]\!] | [\![x]\!]) = 1$, $\kappa([\![Obs^\star(x)]\!] | [\![x]\!]) = 2$, and a rank of 3 to all other observations. Given an initial ranking over worlds, this fixes a ranking over runs. ∎

From a computational perspective, the Markov assumption admits further advantages, particularly if we are interested in modeling the agent's beliefs about propositions. We see from (5) that if the agent cares only about runs that extend the observations seen so far, then the term $\kappa^{\vec{\psi}}([\![w]\!])$ summarizes the influence of the past observations on current and future beliefs. This means that instead of examining an *arbitrarily long* sequence of past observations, the agent can reconstruct its beliefs using its posterior over truth assignments.

The following theorem shows how the agent can update this ranking of assignments when a new observation is made.

Theorem 4.4: *Let $\mathcal{I} = (\mathcal{R}, \kappa, \pi)$ be a Markovian OS. Then*

$$\begin{aligned} &\kappa^{\vec{\psi}\cdot\varphi}([\![w]\!]) = \\ &\quad \kappa([\![Obs^\star(\varphi)]\!] | [\![w]\!]) + \kappa^{\vec{\psi}}([\![w]\!]) - \\ &\quad \min_{\{w:\kappa^{\vec{\psi}}([\![w]\!]) \neq \infty\}} (\kappa([\![Obs^\star(\varphi)]\!] | [\![w]\!]) + \kappa^{\vec{\psi}}([\![w]\!])). \end{aligned}$$

This is the qualitative analogue of standard Bayesian update of a probability distribution using Bayes rule. Since $\kappa^{\vec{\psi}\cdot\varphi}([\![\sigma]\!]) = \min_{w \in [\![\sigma]\!]}(\kappa^{\vec{\psi}\cdot\varphi}([\![w]\!]))$ for $\sigma \in \mathcal{L}$, this theorem shows that all we need to know to compute $\kappa^{\vec{\psi}\cdot\varphi}([\![\sigma]\!])$ is $\kappa^{\vec{\psi}}([\![w]\!])$ for each world w, together with the transition likelihoods. Thus, in many cases, the information an agent needs needs to be able to do revision in this setting, given the Markov assumption, is feasible. Of course, it is still nontrivial to *represent* all this information. But this is precisely the same problem that arises in the probabilistic setting; we would expect the techniques that have proved so successful in the probabilistic setting (for example, Bayesian networks) to be applicable here as well [19].

[6]Note that this example can be captured by a Markovian system if we model the state of the world so that it encodes the recipient's response history as well as her salary.

4.1 Expressive Power of Markovian Systems

One property that immediately follows from the definition of a Markovian OS is the fact that the *order* in which the observations from a sequence are made does not influence the beliefs of the agent; only their presence and quantity do. This suggests the following *exchangability* postulate:

(O2) For any finite observation sequence $\varphi_1, \ldots, \varphi_m$ and for any permutation ρ of $1, \ldots, m$

$$B(\langle \varphi_1, \ldots, \varphi_m \rangle) = B(\langle \varphi_{\rho(1)}, \ldots, \varphi_{\rho(m)} \rangle).$$

It is easy to verify that O2 is sound in Markovian OSs.

Theorem 4.5: *For any Markovian OS $\mathcal{I}$, the revision function $B_{\mathcal{I}}$ induced by $\mathcal{I}$ satisfies O2.*

Unfortunately, O1 and O2 do not suffice to characterize the properties of revision functions in Markovian systems. As we show in the full paper, we can construct revision functions that satisfy both O1 and O2 yet cannot be modeled by a Markovian OS. The question of what conditions are needed to completely characterize Markovian OSs remains open.

5 Credible Observation Models

We have not (yet) put any constraints on the plausibility of observations in different runs. Thus, we can easily construct systems that obey the "opposite" of the success postulate. Consider, for example, a Markovian system where an observation φ is maximally plausible in a world where $\neg\varphi$ holds, and impossible otherwise. That is,

$$\kappa([\![Obs^\star(\varphi)]\!] \mid [\![\alpha]\!]) = \begin{cases} 0 & \text{if } \alpha \models \neg\varphi \\ \infty & \text{if } \alpha \models \varphi. \end{cases}$$

It is easy to verify that in this system, after observing φ, the agent believes $\neg\varphi$.

Of course, this behavior runs counter to our intuition about the role of observations. In this section we examine conditions that attempt to capture the intuition that observations carry useful information about the true state of the world.

We start by considering a very simple condition. We say that an OS is *informative* if an agent is more likely to make accurate observations (ones that are true of the environment state) than inaccurate ones.

Definition 5.1: A Markovian OS $\mathcal{I} = (\mathcal{R}, \kappa, \pi)$ is *informative* if for all $\varphi, \psi \in \mathcal{L}$ and environment states w, if $w \models \varphi \wedge \neg\psi$, then $\kappa([\![Obs^\star(\varphi)]\!] \mid [\![w]\!]) < \kappa([\![Obs^\star(\psi)]\!] \mid [\![w]\!])$.

Informativeness is clearly a nontrivial requirement, and it does seem to go some of the way towards capturing the intuition that observations are usually not misleading. Unfortunately, informative systems need not satisfy even a weak form of success. Consider the OS $\mathcal{I}$ where there are two environment states, p and $\neg p$, such that $\kappa([\![Obs^\star(\top)]\!] \mid [\![p]\!]) = \kappa([\![Obs^\star(\top)]\!] \mid [\![\neg p]\!]) = 0$, $\kappa([\![Obs^\star(p)]\!] \mid [\![p]\!]) = 3$, $\kappa([\![Obs^\star(\neg p)]\!] \mid [\![p]\!]) = 4$, $\kappa([\![Obs^\star(\neg p)]\!] \mid [\![\neg p]\!]) = 1$, $\kappa([\![Obs^\star(p)]\!] \mid [\![\neg p]\!]) = 2$, and $\kappa([\![Obs^\star(\bot)]\!] \mid [\![p]\!]) =$

$\kappa([\![Obs^\star(\bot)]\!] \mid [\![\neg p]\!]) = \infty$. In this system, the only observation the agent is likely to make is the trivial observation $\top$; both p and $\neg p$ are unlikely. $\mathcal{I}$ is informative, since p is more likely to be observed than $\neg p$ if p is true (and $\neg p$ is more likely to be observed than p if $\neg p$ is true). Unfortunately, the agent is still more likely to observe p when p is false than when p is true. Suppose now that the initial ranking is such that both environment states are equally plausible, i.e., $\kappa([\![p]\!]) = \kappa([\![\neg p]\!]) = 0$. It is easy to verify that after m successive observations of p, we have

$$\begin{aligned}\kappa([\![p, Obs(\langle p,\ldots,p\rangle)]\!]) &= 3m\\ \kappa([\![\neg p, Obs(\langle p,\ldots,p\rangle)]\!]) &= 2m.\end{aligned}$$

Thus, $B_{\mathcal{I}}(\langle p,\ldots,p\rangle) = Cn(\neg p)$. Moreover, observing more instances of p only strengthens the belief that p is false.

In our marketing example, this type of situation might arise if we take into account that respondents may be unresponsive (i.e., provide no salary information). Suppose that people with a salary of over \$100,000 are most likely to be unresponsive, while people with a salary of \$90,000 are more likely to report "10" than those that actually have a salary of \$100,000. (The system is informative as long as people with a salary of \$90,000 are more likely to report "9" than "10.") If we take p to denote "10 or more," then this situation is modeled by the system $\mathcal{I}$ above. With each observation "10," our agent (correctly) assesses \$90,000 to be more likely.

Informativeness fails to lead to the recovery of the success postulate because it requires only that we compare the plausibilities of different observations at the same environment state. For an observation φ to provide evidence that increases the agent's degree of belief in φ, we must compare the plausibility of the observation across different states. This suggests the following definition.

Definition 5.2: A Markovian OS $\mathcal{I} = (\mathcal{R}, \kappa, \pi)$ is *credible* if, for all φ and environment states w, v such that $w \models \varphi$ and $v \models \neg\varphi$, we have $\kappa([\![Obs^\star(\varphi)]\!] \mid [\![w]\!]) < \kappa([\![Obs^\star(\varphi)]\!] \mid [\![v]\!])$

Intuitively, credibility says that an observation φ provides stronger evidence for *every* φ-world than it does for *any* $\neg\varphi$-world. In the example above, this requirement is not satisfied, for an observation "10" is most likely to be made in a state where \$90,000 holds.

This requirement allows OSs to satisfy certain weak variants of the success postulate.

(O3) If $\neg\varphi \notin B(\vec{\psi})$, then $\varphi \in B(\vec{\psi}\cdot\varphi)$.

(O4) For all finite observation sequences $\vec{\psi}$ such that $\neg\varphi \notin B(\vec{\psi}\cdot\vec{\psi}')$ for some $\vec{\psi}'$, there is a number n such that $\varphi \in B(\vec{\psi}\cdot\varphi^n)$, where φ^n denotes n repetitions of φ.

Condition **O3** says that if φ is considered plausible (i.e., the agent does not believe $\neg\varphi$), then observing φ suffices to convince the agent that φ is true. Condition **O4** says that if φ is compatible with observing $\vec{\psi}$ (in that there is some sequence of observations that would make φ plausible), then the agent will believe φ after some number of φ observations.

Theorem 5.3: *If $\mathcal{I}$ is credible, then $B_{\mathcal{I}}$ satisfies O3 and O4.*

We can relate informativeness and credibility by requiring an additional property. Suppose that we think of each observation as arising from an experiment that tests the truth or falsity of a particular proposition. Specifically, after experiment E_φ, the agent observes either φ or $\neg\varphi$. In general, whether or not a particular experiment is performed will depend on the state. An OS is *experiment-independent* if the likelihood that a particular experiment is chosen is independent of the state.

Definition 5.4 : The Markovian OS $\mathcal{I} = (\mathcal{R}, \kappa, \pi)$ is *experiment-independent* if, for every observation φ, there is a constant κ_φ, such that $\min(\kappa([\![Obs^\star(\varphi)]\!] \mid [\![w]\!]), \kappa([\![Obs^\star(\neg\varphi)]\!] \mid [\![w]\!])) = \kappa_\varphi$ for all w.

We can think of κ_φ as the likelihood that the experiment for φ is performed (since the experiment will result in observing either φ or $\neg\varphi$). The OS $\mathcal{I}$ described above is not experiment-independent since both observations p and $\neg p$ are less plausible in state p than in state $\neg p$.

While the assumption of experiment-independence does not seem very natural, together with informativeness it implies credibility.

Lemma 5.5: *If $\mathcal{I} = (\mathcal{R}, \kappa, \pi)$ is informative and experiment-independent Markovian OS, then $\mathcal{I}$ is credible.*

Credibility, while allowing for O3 and O4, is still not strong enough to recover the success postulate. For this, we require a yet stronger assumption: we need to assume that all observations are known to be correct; that is, it must be impossible to make an inaccurate observation.

Definition 5.6 : A Markovian OS is *accurate* if $\kappa([\![Obs^\star(\varphi)]\!] \mid [\![w]\!]) = \infty$ whenever $w \not\models \varphi$.

In our example, accuracy requires that when our agent observes, say, "10," the respondent's salary is in fact \$100,000, and the agent is aware of this fact. It is thus impossible to observe two contradictory propositions in any sequence.

Accuracy almost implies informativeness and credibility, but doesn't quite: if observing both φ and $\neg\varphi$ is impossible, the system is accurate but neither informative nor credible. However, as long as $\min(\kappa([\![Obs^\star(\varphi)]\!] \mid [\![w]\!]), \kappa([\![Obs^\star(\neg\varphi)]\!] \mid [\![w]\!])) < \infty$ for all $\varphi \in \mathcal{L}$ and environment states w, then accuracy implies both properties. More significantly, accuracy implies the success postulate.

Theorem 5.7: *If $\mathcal{I}$ is an accurate Markovian OS, then $B_{\mathcal{I}}$ satisfies R2, that is, $\varphi \in B_{\mathcal{I}}(\vec{\psi}\cdot\varphi)$.*

Accuracy is not enough by itself to recover all of the AGM postulates. As we show in the full paper, we need both accuracy and a strong form experiment-independence, which says that all experiments are quite likely (and equally likely) to be performed; that is, $\kappa_\varphi = 0$ for all φ.

The key point here is that, while we can impose conditions on OSs that allow us to recover the full AGM model, it should be clear that these requirements are not always met by naturally-occurring OSs. Notions such as informativeness, credibility and accuracy are often inapplicable in many domains (including our running example of a marketing agent). The framework of OSs (and Markovian OSs in particular) provides a convenient and coherent model for examining the assumptions that hold in a given application domain and determining the precise form a revision function should take.

6 Concluding Remarks

We have described a general ontology for belief revision that allows us to model noisy or unreliable observations and relax the success postulate in a natural way. By imposing the Markov assumption, we obtained OSs that can be easily and naturally described. These give rise to agents whose epistemic state can be encoded in the "usual" way: as a ranking over worlds. Further assumptions about the quality of the observation model allow us to recover the success postulate (and a weaker version of it); this illustrates the general nature of our framework. The emphasis on semantics, as opposed to postulates, has allowed us to readily identify these assumptions and examine their consequences.

There is considerable related work that we survey in detail in a longer version of this paper. Lehmann [26] describes a model where observation sequences are treated as epistemic states in order to deal effectively with iterated revision. Two proposals impact strongly on this paper. Friedman and Halpern [16] use interpreted systems to model both revision and update, and examine the Markov assumption in this context. Boutilier [6, 8] develops a less general model for revision and update (taking the Markov assumption as given) and considers several methods for modeling noisy observations. All of this work (with the exception of [8]) essentially takes the success postulate as a given. Spohn's method of α-conditioning [31], a generalization of the notion of conditioning rankings defined above, was one of the first revision models to explicitly account for strength of evidence. However, α-conditioning does not provide an account of how strength of evidence might be derived. Our model allows us to do this in a natural way, by adapting well-known techniques from probability theory.

Important future research on observation systems includes the incorporation of system dynamics that allows the environment state to change, the development of suitable languages and logics for reasoning with noisy observations, and the syntactic characterization of special cases of OSs (in particular, Markovian OSs). We hope to report on this in future work.

Acknowledgements

The authors are grateful to Daniel Lehmann for comments on a draft of this paper.

References

[1] E. W. Adams. *The Logic of Conditionals*. D.Reidel, 1975.

[2] C. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *J. Sym. Logic*, 50:510–530, 1985.

[3] K. J. Aström. Optimal Control of Markov Decision Processes with Incomplete State Estimation. *J. Math. Analysis and Applications*, 10:174–205, 1965.

[4] C. Boutilier. Normative, subjunctive and autoepistemic defaults: Adopting the Ramsey test. In *Proc. 3rd Inter. Conf. on Principles of Knowledge Representation and Reasoning*, pp. 685–696, 1992.

[5] C. Boutilier. Unifying default reasoning and belief revision in a modal framework. *Artificial Intelligence*, 68:33–85, 1994.

[6] C. Boutilier. Generalized update: Belief change in dynamic settings. In *Proc. 14th Inter. Joint Conf. on AI*, pp. 1550–1556, 1995.

[7] C. Boutilier. Iterated revision and minimal revision of conditional beliefs. *J. Phil. Logic*, 25(3):262–305, 1996.

[8] C. Boutilier. A unified model of qualitative belief change: A dynamical systems perspective. *Artificial Intelligence*, 98(1–2):281–316, 1998.

[9] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.

[10] D. Dubois and H. Prade. Belief change and possibility theory. In Peter Gärdenfors, ed., *Belief Revision*. Cambridge University Press, 1992.

[11] D. Dubois and H. Prade. Belief revision with uncertain inputs in the possibilistic setting. In *Proc. 12th Conf. on Uncertainty in AI*, pp. 236–243, 1996.

[12] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.

[13] N. Friedman. *Modeling Beliefs in Dynamic Systems*. PhD thesis, Stanford University, 1997.

[14] N. Friedman and J. Y. Halpern. Belief revision: A critique. In *Proc. 6th Inter. Conf. on Principles of Knowledge Representation and Reasoning*, pp. 421–431, 1996.

[15] N. Friedman and J. Y. Halpern. A qualitative Markov assumption and its implications for belief change. In *Proc. 12th Conf. on Uncertainty in AI*, pp. 263–273, 1996.

[16] N. Friedman and J. Y. Halpern. Modeling belief in dynamic systems. Part II: revision and update. Submitted for publication. A preliminary version appears in *Proc. 4th Inter. Conf. on Principles of Knowledge Representation and Reasoning*, 1994, pp. 190–201.

[17] N. Friedman and J. Y. Halpern. Modeling beliefs in dynamic systems. part I: Foundations. *Artificial Intelligence*, 95:257–316, 1997.

[18] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, 1988.

[19] M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 84:57–112, 1996.

[20] A. Grove. Two modellings for theory change. *J. Phil. Logic*, 17:157–170, 1988.

[21] J. Y. Halpern and R. Fagin. Modelling knowledge and action in distributed systems. *Distributed Computing*, 3(4):159–179, 1989.

[22] R. C. Jeffrey. *The Logic of Decision*. University of Chicago Press, 1983.

[23] H. Katsuno and A. O. Mendelzon. On the difference between updating a knowledge database and revising it. In *Proc. 2nd Inter. Conf. on Principles of Knowledge Representation and Reasoning*, pp. 387–394, 1991.

[24] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *J. Basic Eng.*, 82:35–45, 1960.

[25] J. G. Kemeny and J. L. Snell. *Finite Markov Chains*. Van Nostrand, 1960.

[26] D. Lehmann. Belief revision, revised. In *Proc. 14th Inter. Joint Conf. on AI*, pp. 1534–1540, 1995.

[27] I. Levi. Iteration of conditionals and the Ramsey test. *Synthese*, 76(1):49–81, 1988.

[28] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[29] H. Rott. Conditionals and theory change: Revisions, expansions, and additions. *Synthese*, 81(1):91–113, 1989.

[30] H. Rott. Belief change using generalized epistemic entrenchment. *J. Logic, Language and Information*, 1(1):45–78, 1992.

[31] W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W.L. Harper and B. Skyrms, eds., *Causation in Decision, Belief Change and Statistics*, vol. 2, pp. 105–134. D. Reidel, 1987.

[32] M. Winslett. *Updating Logical Databases*. Cambridge University Press, 1990.