

**Computer Science Department Technical Report
University of California
Los Angeles, CA 90024-1596**

**REASONING WITH QUALITATIVE PROBABILITIES CAN BE
TRACTABLE**

**M. Goldszmidt
J. Pearl**

**February 1992
CSD-920017**

Reasoning With Qualitative Probabilities Can Be Tractable*

Moisés Goldszmidt[†]

Judea Pearl

< *moises@cs.ucla.edu* > < *judea@cs.ucla.edu* >

Cognitive Systems Laboratory, Computer Science Department,
University of California, Los Angeles, CA 90024

February 20, 1992

Abstract

We recently described a formalism for reasoning with if-then rules that are expressed with different levels of firmness [9]. The formalism interprets these rules as extreme conditional-probability statements, specifying orders of magnitude of disbelief, which impose constraints over possible rankings of worlds. It was shown that, once we compute a priority function Z^+ on the rules, the degree to which a given query is confirmed or denied can be computed in $O(\log n)$ propositional satisfiability tests, where n is the number of rules in the knowledge base. In this paper, we show that computing Z^+ requires $O(n^2 \times \log n)$ satisfiability tests, not an exponential number as was conjectured in [9], which reduces to polynomial complexity in the case of Horn expressions. We also show how imprecise observations can be incorporated in our formalism and how the popular notions of belief-revision and epistemic entrenchment are embodied naturally and tractably.

Key words: Qualitative reasoning, infinitesimal probabilities, belief revision, complexity analysis.

*This work was supported in part by National Science Foundation grant #IRI-88-21444 and State of California MICRO 90-127.

[†]PhD student, supported by an IBM graduate fellowship 1990-1992.

1 Introduction: Infinitesimal Probabilities, Rankings and Common Sense Reasoning

The uncertainty encountered in common sense reasoning fluctuates over an extremely wide range. For example, the probability that the new book on my desk is about astrology is less than one in a million. If however, I spot a Zodiac sign on page 1, the probability becomes close to 1, say 0.999. Intelligent agents are expected to reason with such rare eventualities and to produce explanations and actions whenever these occur. Given this wide range of uncertainty and the fact that the majority of everyday decisions involve relatively low payoffs, the full precision of probability calculus may not be necessary, and an order-of-magnitude approximation may be sufficient. Thus, instead of measuring probabilities on a scale from zero to one, we can imagine projecting probability measures onto a quantized logarithmic scale and then treating beliefs that map onto two different quanta as being of different orders of magnitude.

This method of approximation gives rise to a semi-qualitative calculus of uncertainty, one in which degrees of (dis)belief are ranked by non-negative integers (corresponding perhaps to linguistic quantifiers such as “believable,” “unlikely,” “very rare”) still capable of accounting for retraction and restoration of beliefs by Bayesian conditionalization. The origin of this approximation can be traced back to Ernest Adams [1], who developed a logic of conditionals based on infinitesimal probabilities, and to the Ordinal Condition Functions of Spohn [18]. Potential applications in nonmonotonic reasoning were noted in [12, 13] and further developed in [11, 7, 15, 8, 9].

A simple way of viewing infinitesimal probabilities is to consider an ordinary probability function P defined over a set Ω of possible worlds (or states of the world) ω and to imagine that the probability $P(\omega)$ is a polynomial function of some small parameter ε , for example $1 - c_1\varepsilon$ or $\varepsilon^2 - c_2\varepsilon^4$. Accordingly, the probabilities assigned to any subset of Ω represented by a logical formula φ , as well as all conditional probabilities $P(\psi|\varphi)$, will be rational functions of ε . We define the ranking function $\kappa(\psi|\varphi)$ ¹ as the power of the most significant ε -term in the expansion of $P(\psi|\varphi)$. In other words, $\kappa(\psi|\varphi) = n$ iff $P(\psi|\varphi)$ has the same order of magnitude as ε^n .

The following properties of ranking functions (left-hand side below) reflect, on a logarithmic scale, the usual properties of probability functions (right-hand side), with “min” replacing addition, and addition replacing multiplication:

$$\kappa(\varphi) = \min_{\omega \models \varphi} \kappa(\varphi) \quad : \quad P(\varphi) = \sum_{\omega \models \varphi} P(\varphi) \quad (1)$$

$$\kappa(\varphi) = 0 \text{ or } \kappa(\neg\varphi) = 0 \text{ or both} \quad : \quad P(\varphi) + P(\neg\varphi) = 1 \quad (2)$$

$$\kappa(\psi \wedge \varphi) = \kappa(\psi|\varphi) + \kappa(\varphi) \quad : \quad P(\psi \wedge \varphi) = P(\psi|\varphi)P(\varphi) \quad (3)$$

¹Spohn [18] called this function a “non-probabilistic” Ordinal Condition Function.

Parameterizing a probability measure by ε and extracting the lowest exponent of ε as the measure of (dis)belief mirrors the process by which people abstract qualitative beliefs from numerical probabilities and accept them as tentative truths. For example, we can make the following correspondence between linguistic quantifiers and ε^n :

$P(\phi) = \varepsilon^0$	ϕ is believable, $\neg\phi$ is believable	$\kappa(\phi) = 0$
$P(\phi) = \varepsilon^1$	ϕ is unlikely, $\neg\phi$ is believed	$\kappa(\phi) = 1$
$P(\phi) = \varepsilon^2$	ϕ is very unlikely, $\neg\phi$ is strongly believed	$\kappa(\phi) = 2$
$P(\phi) = \varepsilon^3$	ϕ is extremely unlikely, $\neg\phi$ is very strongly believed	$\kappa(\phi) = 3$
\vdots	\vdots	\vdots

These approximations yield a probabilistically sound calculus, employing integer addition, for manipulating the orders of magnitude of disbelief. It is governed by the following principles:

1. Each world is ranked by a non-negative integer κ representing the degree of surprise associated with finding such a world.
2. Each well-formed formula (wff) is given the rank of the world with the lowest κ (most normal world) that satisfies that formula.
3. Given a collection of facts ϕ , we say that σ follows from ϕ with strength δ if $\kappa(\psi|\phi) > \delta$, or, equivalently, if the κ rank of $\phi \wedge \neg\sigma$ is at least $\delta + 1$ degrees above that of $\phi \wedge \sigma$.²

The basic κ ranking system, as described in Spohn [18], requires the specification of a complete ranking function before reasoning can commence. In other words, the knowledge base must be sufficiently rich to define the κ associated with every world ω . Unfortunately, in practice such specification is not readily available. For example, we might be given the information that “birds normally fly” (written $\kappa(\neg f|b) > 0$) and no information whatsoever about the flying habits of red birds or non-birds. We still would like to conclude that red birds normally fly, even though the information given is not sufficient for defining a complete ranking function. In order to draw plausible conclusions from such fragmentary pieces of information, we require additional inferential machinery that should accomplish two functions: First, it should enrich the specification of the ranking function with the needed information and, second, it should operate directly on the specification sentences in the knowledge base, rather than on the rankings of worlds (which are too numerous to list). Such machinery is provided by a formalism called system- Z^+ [9] which accepts knowledge in the form of quantified if-then rules (e.g., “birds fly (with strength δ)”) and computes the plausibility of any given query (e.g., “Tim, being a red-bird, flies (with degree δ')”) by syntactic manipulation of these rules.

To accomplish these functions, system- Z^+ incorporates two principles in addition to those given above:

²In probabilistic terms, $P(\sigma|\phi) \geq 1 - c\varepsilon^{(\delta+1)}$.

4. Each input rule “if φ then ψ (with strength δ),” written $\varphi \xrightarrow{\delta} \psi$, is interpreted as a constraint on the ranking κ , forcing every world in $\varphi \wedge \neg\psi$ to rank at least $\delta + 1$ degrees above the most normal world in φ , that is, $\kappa(\psi|\varphi) > \delta$.
5. Out of all rankings satisfying the constraints above, we adopt the (unique) ranking κ^+ that assigns each world the lowest possible (most normal) rank.

The first contribution of this paper is to improve the inference process of system- Z^+ and establish its tractability. A key step in the procedures developed in [9] was the computation of a priority ranking Z^+ on the rules in the knowledge base, which was conjectured to be intractable. In Section 3 (after some preliminary definitions in Section 2), we present a procedure for computing Z^+ that requires a polynomial number of propositional satisfiability tests and hence is tractable in applications permitting restricted languages, such as Horn expressions, network theories, or acyclic databases.

The second contribution of this paper is to equip system- Z^+ with the capability to reason with *soft* evidence or imprecise observations (Section 4). Such a capability is important when we wish to assess the plausibility of σ (using Principle 3 above) but the context ϕ is not given with absolute certainty. In other words, there is some vague testimony supporting ϕ but that testimony is undisclosed (or cannot be articulated using the basic propositions in our language, e.g., testimony of the senses); all that can be ascertained is a summary of that testimony saying that “ ϕ is supported to a degree n .” We propose two different strategies for computing a new ranking κ' from an initial one κ , given soft evidential report supporting a wff ϕ . The first strategy, named J-conditionalization, is based on *Jeffrey’s Rule of Conditioning* [14]. It interprets the report as specifying that “all things considered,” the new degree of disbelief for $\neg\phi$ should be $\kappa'(\neg\phi) = n$. The second strategy, named L-conditionalization, is based on the *virtual evidence* proposal described in [13]. It interprets the report as specifying the desired *shift* in the degree of belief in ϕ , as warranted by that report alone and “nothing else considered.” We show that L-conditionalization has roughly the same complexity as ordinary conditionalization, and then we relate our formalism to the theory of belief revision in [2]. Finally, Section 5 summarizes the main results.

2 Preliminary Definitions: The Ranking κ^+

We start with a set of rules $\Delta = \{r_i \mid r_i = \varphi_i \xrightarrow{\delta_i} \psi_i, 1 \leq i \leq n\}$, where φ_i and ψ_i are propositional formulas over a finite language of atomic propositions, “ \rightarrow ” denotes a new connective, and δ_i is a non-negative integer. A truth valuation ω of the atomic propositions in the language will be called a *world*. The satisfaction of a wff φ by a world ω is defined as usual and denoted by $\omega \models \varphi$. ω is a *model* for φ if ω satisfies φ . Let Ω stand for the set of possible worlds. *Ranking functions* are defined as follows:

Definition 2.1 (Rankings) A ranking function κ is an assignment of non-negative integers to the elements in Ω , such that $\kappa(\omega) = 0$ for at least one world $\omega \in \Omega$. We extend this definition to induce rankings on wffs:³

$$\kappa(\varphi) = \begin{cases} \min_{\omega \models \varphi} \kappa(\omega) & \text{if } \varphi \text{ is satisfiable,} \\ \infty & \text{otherwise.} \end{cases} \quad (4)$$

Similarly, given two wffs φ and ψ such that φ is satisfiable, we define the conditional ranking $\kappa(\psi|\varphi)$ as $\kappa(\psi|\varphi) = \kappa(\psi \wedge \varphi) - \kappa(\varphi)$.

□

Definition 2.2 (Consistency) A ranking κ is said to be **admissible** relative to a given Δ , iff

$$\min_{\omega \models \varphi_i \wedge \psi_i} \kappa(\omega) + \delta_i < \min_{\omega \models \varphi_i \wedge \neg \psi_i} \kappa(\omega) \quad (5)$$

(equivalently $\kappa(\varphi_i \wedge \psi_i) + \delta_i < \kappa(\varphi_i \wedge \neg \psi_i)$, or $\kappa(\neg \psi_i | \varphi_i) > 0$) for every rule $\varphi_i \xrightarrow{\delta_i} \psi_i \in \Delta$. A set Δ is **consistent** iff there exists an admissible ranking κ relative to Δ .

□

Consistency can be decided in $O(|\Delta|^2)$ satisfiability tests, and it is independent of the δ -values assigned to the rules in Δ [9]. Eq. 5 echoes the usual interpretation of *defaults* rules [17], according to which ψ holds in all *minimal* models for φ . In our case, minimality is reflected in having the lowest rank, that is the highest possible likelihood. If we say that ω *falsifies* or *violates* a rule $\varphi \xrightarrow{\delta} \psi$ whenever $\omega \models \varphi \wedge \neg \psi$, the parameter δ can be interpreted as the minimal degree of surprise (or abnormality) associated with finding the rule $\varphi \xrightarrow{\delta} \psi$ violated, given that we know φ . In probabilistic terms, consistency guarantees that for every $\varepsilon > 0$, there exists a probability distribution P such that if $\varphi_i \xrightarrow{\delta_i} \psi_i \in \Delta$, then $P(\psi_i | \varphi_i) \geq 1 - c\varepsilon_i^{\delta_i}$.

The distinguished ranking κ^+ , defined below, assigns to each world the lowest possible rank permitted by the admissibility constraints of Eq. 5 (Def. 2.2). Such an assignment reflects the assumption that, unless we are forced to do otherwise, each world is considered as normal (likely) as possible.

Definition 2.3 (The ranking κ^+) Let $\Delta = \{r_i \mid r_i = \varphi_i \xrightarrow{\delta_i} \psi_i\}$ be a consistent set of rules. κ^+ is defined as follows:

$$\kappa^+(\omega) = \begin{cases} 0 & \text{if } \omega \text{ does not falsify any rule in } \Delta, \\ \max_{\omega \models \varphi_i \wedge \neg \psi_i} [Z^+(r_i)] + 1 & \text{otherwise,} \end{cases} \quad (6)$$

³For notational convenience we will use the same symbol κ to denote the ranking on wffs.

where $Z^+(r_i)$ is a *priority* ranking on rules, defined by

$$Z^+(r_i) = \min_{\omega \models \varphi_i \wedge \psi_i} [\kappa^+(\omega)] + \delta_i. \quad (7)$$

□

In [9], we show that Eqs. 6 and 7 define a unique admissible ranking function κ^+ that is minimal in the following sense: Any other admissible ranking function must assign a higher ranking to at least one world and a lower ranking to none.

3 Plausible Conclusions: Computing the Z^+ -rank

Given a set Δ , each admissible ranking κ induces a consequence relation \vdash_{κ} , where $\phi \vdash_{\kappa} \sigma$ iff $\kappa(\sigma \wedge \phi) < \kappa(\neg\sigma \wedge \phi)$. A straightforward way to declare σ as a plausible conclusion of Δ given ϕ would be to require $\phi \vdash_{\kappa} \sigma$ in all κ admissible with Δ . This leads to an entailment relation called ε -semantics [13], 0-entailment [15], and r-entailment [11], which is recognized as being too conservative. The approach we take here is to select a distinguished admissible ranking, in our case κ^+ , and declare σ as a plausible conclusion of Δ given ϕ , written $\phi \vdash_{\kappa^+} \sigma$, iff $\kappa^+(\phi \wedge \sigma) < \kappa^+(\phi \wedge \neg\sigma)$.⁴ According to Eq. 6 in Def. 2.3, both κ^+ and \vdash_{κ^+} can be computed effectively once the priority ranking Z^+ on rules is known. We next present a procedure for computing Z^+ , which is identical to the one presented in [9] save for the crucial computation of Eq. 8 (Step 3(b)). Whereas in [9] this computation was thought to require an exponential search over worlds, we now show that it can be accomplished in $O(|\Delta| \times \log |\Delta|)$ satisfiability tests. Some of the steps in Procedure **Z.rank** depend upon the notion of *toleration*. A rule $\phi \xrightarrow{\delta} \sigma$ is *tolerated* by Δ if the wff $\phi \wedge \sigma \wedge \varphi_i \supset \psi_i$ is satisfiable (where i ranges over all rules in Δ).⁵

Procedure **Z.rank**

Input: A consistent knowledge base Δ . **Output:** Z^+ -ranking on rules.

1. Let Δ_0 be the set of rules tolerated by Δ , and let $\mathcal{R}Z^+$ be an empty set.
2. For each rule $r_i = \varphi_i \xrightarrow{\delta_i} \psi_i \in \Delta_0$, do: set $Z(r_i) = \delta_i$ and $\mathcal{R}Z^+ = \mathcal{R}Z^+ \cup \{r_i\}$.
3. While $\mathcal{R}Z^+ \neq \Delta$, do:

⁴If we are concerned with the strength δ with which the conclusion is endorsed, then $\phi \vdash_{\kappa^+}^{\delta} \sigma$ iff $\kappa^+(\phi \wedge \sigma) + \delta < \kappa^+(\phi \wedge \neg\sigma)$.

⁵The notion of toleration is also crucial for deciding consistency: Δ is consistent iff there is a tolerated default rule in every nonempty subset Δ' of Δ (Theorem 1, [9]).

- (a) Let Δ^+ be the set of rules in $\Delta' = \Delta - \mathcal{RZ}^+$ tolerated by Δ'
- (b) For each $r : \phi \xrightarrow{\delta} \sigma \in \Delta^+$, let Ω_r denote the set of models for $\phi \wedge \sigma$ that do not violate any rule in Δ' : compute

$$Z(r) = \min_{\omega_r \in \Omega_r} [\kappa(\omega_r)] + \delta \quad (8)$$

where

$$\kappa(\omega_r) = \max_{r_j \in \mathcal{RZ}^+} \{Z(r_j) \mid \omega_r \models \varphi_j \wedge \neg \psi_j\} + 1 \quad (9)$$

and $r_j : \varphi_j \xrightarrow{\delta_j} \psi_j \in \mathcal{RZ}^+$.⁶

- (c) Let r^* be a rule in Ω having the lowest Z ; set $\mathcal{RZ}^+ = \mathcal{RZ}^+ \cup \{r^*\}$.

End Procedure

Theorem 3.1 about the correctness of Procedure `Z_rank` is taken from [9]. Lemmas 3.2 and 3.3 and Theorem 3.4 prove the polynomial complexity of `Z_rank`.

Theorem 3.1 *The function Z computed by `Z_rank` complies with Def. 2.3, that is $Z = Z^+$.*

Lemma 3.2 *Let $\Delta = \{r_i \mid r_i = \varphi_i \xrightarrow{\delta_i} \psi_i\}$ be a consistent set where the rules are sorted in nondecreasing order according to priorities $Z(r_i)$. Let $\kappa(\omega)$ be defined as in Eq. 6:*

$$\kappa(\omega) = \begin{cases} 0 & \text{if } \omega \text{ does not falsify any rule in } \Delta, \\ \max_{\omega \models \varphi_i \wedge \neg \psi_i} [Z(r_i)] + 1 & \text{otherwise.} \end{cases} \quad (10)$$

Then, for any wff ϕ , $\kappa(\phi)$ can be computed in $O(\log |\Delta|)$ propositional satisfiability tests.

Proof: The idea is to perform a binary search on Δ to find the lowest $Z(r)$ such that there is a model for ϕ that does not violate any rule r' with priority $Z(r') \geq Z(r)$. We first divide Δ into two roughly equal sections: top-half (r_{mid} to r_{high}) and bottom-half (r_{low} to r_{mid}). Then we examine the top-half: If the wff $\alpha = \phi \wedge_{j=mid}^{j=n} \varphi_j \supset \psi_j$ is satisfiable, then there exists a model for ϕ that does not violate any rule in this top-half. It follows that $Z(r_{mid}) + 1$ is an upper bound on the value of $\kappa(\phi)$, and the binary search is continued iteratively in the bottom-half. If, on the other hand, α is not satisfiable, then the maximum $Z(r_i)$ for any model for ϕ must be in the top-half, and the search is continued there. Eventually, the set in which the search is conducted is reduced to one rule, and we can determine the value of $\kappa(\phi)$ with one more satisfiability test. \square

⁶Note that Eqs. 8 and 9 correspond to Eqs. 7 and 6 in Def. 2.3.

Lemma 3.3 *The value of $Z(\phi \xrightarrow{\delta} \sigma)$ in Eq. 8 can be computed in $O(\log |\mathcal{R}Z^+|)$ satisfiability tests.*

Proof: Let Δ' in Step 3(a) be equal to $\{\varphi_i \xrightarrow{\delta_i} \psi_i\}$, and let the wff α be equal to $\sigma \wedge \phi \wedge_i \varphi_i \supset \psi_i$ where i ranges over all the rules in Δ' . Note that since any world ω_r in Ω_r is a model for $\sigma \wedge \phi$ and does not violate any rule in Δ' , it follows that $\omega_r \in \Omega_r$ iff $\omega_r \models \alpha$. Then, since $\kappa(\alpha) = \min_{\omega_r \in \Omega_r} \kappa(\omega_r)$, $Z(\phi \xrightarrow{\delta} \sigma)$ must be equal to $\kappa(\alpha) + 1 + \delta$. Thus, once Δ' is sorted, by Lemma 3.2, $\kappa(\alpha)$ can be computed in $O(\log |\mathcal{R}Z^+|)$ satisfiability tests, which proves Lemma 3.3. \square

Theorem 3.4 *Given a consistent Δ , the computation of the ranking Z^+ requires $O(|\Delta|^2 \times \log |\Delta|)$ satisfiability tests.*

Proof: Step 1 requires at most $|\Delta|$ satisfiability tests and is performed once, while Step 2 takes at most $|\Delta|$ data assignments. Step 3(a) again requires $O(|\Delta|)$ satisfiability tests. Computing Eq. 8 in Step 3(b) can be done in $O(\log |\mathcal{R}Z^+|)$ satisfiability tests according to Lemma 3.3,⁷ and since it will be executed at most $O(|\Delta|)$ times, it requires a total of $O(|\Delta| \times \log |\Delta|)$ satisfiability tests. Step 3(c) is a minimum search which can be done in conjunction with the computation of Eq. 8, since we only need to keep the minimum of such values. It involves $|\Omega|$ data comparisons. Loop 3 is performed at most $|\Delta| - |\Delta_0|$ times, hence the whole computation of the priorities Z^+ on rules requires a total of $O(|\Delta|^2 \times \log |\Delta|)$ satisfiability tests. \square

Once Z^+ is known, determining the strength δ with which an arbitrary query σ is confirmed, given the information ϕ , requires $O(\log |\Delta|)$ satisfiability tests: First $\kappa^+(\phi \wedge \sigma)$ and $\kappa^+(\phi \wedge \neg\sigma)$ are computed, using a binary search as in Lemma 3.2. Then, these two values are compared and the difference is equated with the strength δ . Clearly, if the rules in Δ are of Horn form, computing the priority ranking Z^+ and deciding the plausibility of queries $(\phi \xrightarrow{\delta} \sigma)$ can be done in polynomial time [4].

4 Belief Change, Soft Evidence, and Imprecise Observations

So far, a query $\phi \xrightarrow{\delta} \sigma$ was defined as a pair of Boolean formulae (ϕ, σ) , where ϕ (the *context*) stands for the set of observations at hand and σ (the *target*) stands for the conclusion whose

⁷Note that we need $\mathcal{R}Z^+$ to be sorted, nondecreasingly, with respect to the priorities Z . This requires that the initial values inserted in $\mathcal{R}Z^+$ in Step 2 of Procedure `Z_rank` be sorted — $O(|\Delta_0|^2)$ data comparisons — and that the new Z -value in Step 3(c) be inserted in the right place — $O(|\mathcal{R}Z^+|)$ data comparisons. We are assuming that the cost of each of these operations is much less than that of a satisfiability test.

belief we wish to confirm, deny, or assess. A query (ϕ, σ) would be answered in the affirmative if σ was found to hold in all minimally ranked models of ϕ , and the degree of belief in σ would be given by $\kappa(\neg\sigma|\phi) - \kappa(\sigma|\phi)$.

In many cases, however, the queries we wish to answer cannot be cast in such a format, because our set of observations is not precise enough to be articulated as a crisp Boolean formula. For example, assume that we are throwing a formal party and our friends Mary and John are invited. However, judging from their previous behavior, we believe that “if Mary goes to the party, John will stay home (with strength δ),” written $M \xrightarrow{\delta} \neg J$. Now assume that we have a strong hunch (with degree K) that Mary will go to the party (perhaps because she is extremely well dressed and is not consulting the movie section in the *Times*) and we wish to inquire whether John will stay home. It would be inappropriate to query the system with the pair $(M, \neg J)$, because the context M has not been established beyond doubt. The difference could be critical if we have arguments against “John staying home,” for example, that he was seen renting a tuxedo. A flexible system should allow the user to assign a degree of belief to each observational proposition in the context ϕ and proceed with analyzing their rational consequences. Thus, a query should consist of a tuple like $(\phi_1, K_1, \phi_2, K_2, \dots, \phi_m, K_m : \sigma)$, where each K_i measures the degree to which the contextual proposition ϕ_i is supported by evidence.

At first glance it might seem that such facility is automatically provided by system- Z^+ , through the use of variable-strength rules. For example, to express the fact that Mary is believed to be going to the party, we can perhaps use a *dummy* rule $Obs_1 \xrightarrow{K} M$ (stating that if Mary meets the set of observations Obs_1 then Mary is believed to be going to the party) and then add the proposition Obs_1 to the context part of the query, to indicate that Obs_1 has taken place.

This proposal has several shortcomings, however. First, the net impact of our new rule $Obs_1 \xrightarrow{K} M$ would be sensitive to previously collected information about Mary’s intentions (say she has bought a plane ticket) that we may wish to suppress. In other words, we often wish to express the assessment that, **all things considered**, Mary’s going to the party is believed to degree K .

Second, in many systems it is convenient to treat if-then rules as a stable part of our knowledge, unperturbed by observations made about a particular individual or in any specific set of circumstances. This permits us to compile rules into a structure that allows efficient query processing. Adding query-induced rules to the knowledge base will neutralize this facility.

Finally, rules and observations combine differently: The latter should accumulate the former do not. For example, if we have two rules $a \xrightarrow{\delta_1} c$ and $b \xrightarrow{\delta_2} c$ and we observe a and b , system- Z^+ would believe c to a degree $\max(\delta_1, \delta_2)$. However, if a and b provide two independent reasons for believing c , the two observations together should endow c with a belief that is stronger than any one component in isolation. To incorporate such cumulative pooling of evidence, we must encode the assumption that a and b are conditionally independent (given

c), which is not automatically embodied in system- Z^+ .⁸

To avoid these complications, the method we propose treats imprecise observations by invoking specialized conditioning operators, unconstrained by a rule’s semantics. We distinguish between two types of evidential reports:

1. Type-J: “All things considered,” our current belief in ϕ should **become** J .
2. Type-L: “Nothing else considered,” our current belief in ϕ should **shift by** L .

4.1 Type-J: All Things Considered

Let ϕ be the wff representing the event whose belief we wish to update so that $\kappa'(\neg\phi) = J$ (and, consequently, $\kappa'(\phi) = 0$). In order to compute $\kappa'(\psi)$ for every wff ψ , we rely upon *Jeffrey’s Rule of Conditioning* [14]. Jeffrey’s rule is based on the assumption that while the observation changed the agent’s degree of belief in ϕ and in certain other proposition, it did not change the **conditional degree of belief** in any propositions on the evidence ϕ or on the evidence $\neg\phi$ [14]. Thus, letting P' denote the agent’s probability distribution after the report on the value of $P'(\phi)$ is incorporated, and, using P to denote the agent’s probability distribution prior to this report, we have⁹

$$P'(\psi|\phi) = P(\psi|\phi) \quad \text{and} \quad P'(\psi|\neg\phi) = P(\psi|\neg\phi), \quad (11)$$

which leads to Jeffrey’s rule,

$$P'(\psi) = P(\psi|\phi)P'(\phi) + P(\psi|\neg\phi)P'(\neg\phi). \quad (12)$$

Translated into the language of rankings (using Eqs. 1–3), Eq. 12 yields

$$\kappa'(\psi) = \min[\kappa(\psi|\phi) + \kappa'(\phi); \kappa(\psi|\neg\phi) + \kappa'(\neg\phi)], \quad (13)$$

which offers a convenient way of computing $\kappa'(\psi)$ once we specify $\kappa'(\phi) = 0$ and $\kappa'(\neg\phi) = J$. Eq. 13 assumes the an especially attractive form when computing the κ' of a world ω :

$$\kappa'(\omega) = \begin{cases} \kappa(\omega|\phi) + \kappa'(\phi) & \text{if } \omega \models \phi \\ \kappa(\omega|\neg\phi) + \kappa'(\neg\phi) & \text{if } \omega \models \neg\phi \end{cases} \quad (14)$$

⁸The assumptions of conditional independence among converging rules is embodied in the formalism of Maximum Entropy [8].

⁹Eq. 11 is known as the *J-condition* [14].

Eq. 14 corresponds exactly to the α -conditionalization proposed in Spohn [18] (Def. 6, p. 117), with $\alpha = J$. If $\kappa'(\neg\phi) = \infty$, this process is equivalent to ordinary Bayesian conditionalization, since $k'(\omega) = k(\omega|\phi)$ if $\omega \models \phi$ and $\kappa'(\omega) = \infty$ otherwise. Note, however, that in general this conditionalization is not commutative; if ϕ_1 and ϕ_2 are mutually dependent (i.e., $\kappa(\phi_2|\phi_1) \neq \kappa(\phi_2)$),¹⁰ the order in which we establish $\kappa(\neg\phi_1) = J_1$ and $\kappa(\neg\phi_2) = J_2$ might make a difference in our final belief state, represented by the ranking k'' .¹¹

4.2 Type-L Reports: Nothing Else Considered

L-conditionalization is appropriate for evidential reports of the type “a new evidence was obtained which, by its own merit, would *support* ϕ to degree L .” Unlike J-conditionalization, the degree L now specifies *changes* in the belief of ϕ , not the absolute value of the final belief in ϕ . As in the case of type-J reports, we assume that in naming ϕ as the direct beneficiary of the evidence, the intent is to convey the assumption of conditional independence, as formulated in Eq. 12. Next, we assume that the degree of support L characterizes the likelihood-ratio $\lambda(\phi)$ associated with some undisclosed observation Obs , as is done in the method of *virtual conditionalization* [13]:

$$\lambda(\phi) = \frac{P(Obs|\phi)}{P(Obs|\neg\phi)}, \quad (15)$$

which governs the updates via the product rule

$$\frac{P'(\phi)}{P'(\neg\phi)} = \frac{\lambda(\phi)P(\phi)}{P(\neg\phi)}. \quad (16)$$

Translated into the language of rankings, this assumption yields

$$\kappa'(\phi) - \kappa'(\neg\phi) = \kappa(\phi) - \kappa(\neg\phi) - L \quad (17)$$

and, since either $\kappa'(\phi)$ or $\kappa'(\neg\phi)$ must be zero, we obtain

$$\kappa'(\phi) = \max[0; \kappa(\phi) - \kappa(\neg\phi) - L], \quad (18)$$

$$\kappa'(\neg\phi) = \max[0; \kappa(\neg\phi) - \kappa(\phi) + L]. \quad (19)$$

¹⁰This condition mirrors probabilistic dependence, i.e., $P(\phi_2|\phi_1) \neq P(\phi_2)$.

¹¹Spohn ([18], p. 118) has acknowledged the desirability of commutativity in evidence pooling but has not stressed that α -conditionalization commutes only in a very narrow set of circumstances (partially specified by his Theorem 4). These circumstances require that successive pieces of evidence support only propositions that are relatively independent — the truth of one proposition should not imply a belief in another. Shenoy [16] has corrected this deficiency by devising a commutative combination rule which behaves similar to L-conditioning.

We see that the effect of L-conditionalization is to shift the degree of disbelief difference between ϕ and $\neg\phi$ by the specified amount L . Once $\kappa'(\phi)$ is known, we can use Jeffrey's rule (Eq. 13) to compute the $\kappa'(\psi)$ for an arbitrary wff ψ

$$\kappa'(\psi) = \begin{cases} \min[\kappa(\psi|\phi) + \kappa(\phi) - \kappa(\neg\phi) - L; \kappa(\psi|\neg\phi)] & \text{if } \kappa(\neg\phi) - \kappa(\phi) < L, \\ \min[\kappa(\psi|\phi); \kappa(\psi|\neg\phi) + \kappa(\neg\phi) + L - \kappa(\phi)] & \text{if } \kappa(\neg\phi) - \kappa(\phi) > L, \\ \min[\kappa(\psi|\phi); \kappa(\psi|\neg\phi)] & \text{if } \kappa(\neg\phi) - \kappa(\phi) = L. \end{cases} \quad (20)$$

This expression takes the following form for $\kappa'(\omega)$:

$$\kappa'(\omega) = \begin{cases} \kappa(\omega|\phi) + \max[0; \kappa(\phi) - \kappa(\neg\phi) - L] & \text{if } \omega \models \phi, \\ \kappa(\omega|\neg\phi) + \max[0; \kappa(\neg\phi) - \kappa(\phi) + L] & \text{if } \omega \models \neg\phi. \end{cases} \quad (21)$$

As in J-conditionalization, if $L = \infty$ then $\kappa'(\omega) = \kappa(\omega|\phi)$. For the general case, we can see that the effect of L-conditionalization is to shift *downward* the κ of all worlds that are models of the supported proposition ϕ relative to the κ of all worlds that are not models for ϕ . However, unlike J-conditionalization, the net relative shift is constant and is equal to L , independent of the initial value of $\kappa(\phi)$. It is easy to verify that L-conditionalization is commutative (as is its probabilistic counterpart, see Eq. 16), and hence it permits a recursive implementation in the case of multiple evidence.

We can illustrate these updating schemes through the party example consisting of the single rule $r_m : M \xrightarrow{4} \neg J$ (“if Mary goes to the party, then John will not go”). A trivial application of Procedure `Z_rank` yields $Z^+(r_m) = 4$, and using Eqs. 4 and 6 we find $\kappa(x) = 0$ for every proposition x , except $x = J \wedge M$, for which we have $\kappa^+(M \wedge J) = 5$. This means that we have no reason to believe that either Mary or John will go to the party, but we are pretty sure that both of them will not show up. Now suppose we see that Mary is very well dressed, and this observation makes our belief in M increase to 3, that is, $\kappa^+(\neg M) = 3$. As a consequence, our belief in John staying home also increases to 3 since, using either J-conditionalization or L-conditionalization, $\kappa^+(J) = 3$. Next, suppose that someone tells us that he has a strong hunch that John plans to show up for the party, but he fails to tell us why. There are two ways in which this report can influence our beliefs. The natural way would be to assume that our informer has not seen Mary's dress, and might not even be aware of John and Mary's relationship — hence we assess the impact of his report in isolation and say that whatever the value of our current belief in John going, it should increase by 3 increments, or $L = 3$. Following Eq. 20, $\kappa^{+''}(J)$ and $\kappa^{+''}(\neg M)$ will both be equal to 0, and we are back to the initial uncertainty about John or Mary going to the party, except that our disbelief in both Mary and John being at the party has decreased to $\kappa^{+''}(M \wedge J) = 2$. The second way would be to assume that our informer is omniscient and already has taken into consideration all we know about John and Mary. He means for us to revise our rankings so that the final belief in “John going” will be fixed at $\kappa^{+''}(\neg J) = 3$. With this interpretation, we J-condition $\kappa^{+''}$ on the proposition $\phi = \neg J$ and obtain $\kappa^{+'''}(M) = 3$, concluding that it is Mary who will not show up to the party after all.

4.3 Complexity Analysis

From Eq. 13 we see that $\kappa'(\psi)$ can be computed from $\kappa(\psi|\phi)$ and $\kappa(\psi|\neg\phi)$, which, assuming we have Z^+ , requires a logarithmic number of propositional satisfiability tests (see Section 3). L-conditionalization can follow a similar route, as depicted in Eq. 20.

Special precautions must be taken when simultaneous, multiple pieces of evidence become available. First, J-conditionalization is not commutative, hence we cannot simply compute κ' by J-conditioning on ϕ_1 and then J-conditioning κ' on ϕ_2 to get κ'' . We must J-condition simultaneously on ϕ_1 and ϕ_2 with their respective J-levels, say J_1 and J_2 . Worse yet, an auxiliary effort must be expended to compute the J-level of each combination of ϕ 's, in our case $\phi_1 \wedge \phi_2$, $\phi_1 \wedge \neg\phi_2$, etc. This is no doubt a hopeless computation when the number of observations is large.

L-conditionalization, by virtue of its commutativity, enjoys the benefits of recursive computations. Let e_1 and e_2 be two (undisclosed) pieces of evidence supporting ϕ_1 (with strength L_1) and ϕ_2 (with strength L_2), respectively. We first L-condition κ on ϕ_1 and calculate $\kappa'(\phi_1)$ and $\kappa'(\phi_2)$ using Eq. 19 and Eq. 20, respectively. Applying Eq. 20 this time to $\kappa'(\psi \wedge \phi_2)$, we calculate $\kappa'(\psi|\phi_2)$. Second, we L-condition κ' on ϕ_2 , compute $\kappa''(\phi_2)$ using Eq. 19, and, finally, using $\kappa'(\psi|\phi_2)$ and $\kappa''(\phi_2)$ in Eq. 20 obtain $\kappa''(\psi)$.¹² Note that, although each of these calculations requires only $O(\log |\Delta|)$ satisfiability tests, this computation is effective only when we have a well designated target hypothesis ψ to estimate. The computation must be repeated each time we change the target hypothesis, even when the context remains unaltered. This is so because we no longer have a facility for encoding a complete description of κ' , as we had for κ (using the Z^+ -function). Thus, the encoding for κ' may not be as *economical* as that for κ (the number of worlds is astronomical), unless we manage to find dummy rules that emulate the constraints imposed on ϕ_1 by the (undisclosed) observation. Such dummy rules must enforce the conditional independence constraints embedded in Eq. 12, without violating the admissibility constraints (Eq. 5) in Δ (see [10]).

4.4 Relation to the AGM Theory of Revision

Alchourrón, Gärdenfors and Makinson (AGM) have advanced a set of postulates which define a minimal set of desirable properties for rational belief revision [2]. These postulates have become a standard, against which (new) proposals are being tested. However, as was pointed out by Boutilier [3] and others, the AGM postulates do not provide a calculus with which one can either reason about or, realize the revision process.

Spohn [18] has shown that if one associates the set of beliefs with those propositions ψ for which $\kappa(\neg\psi) > 0$, then the process of belief revision affected by α -conditioning (or J-conditioning in our terminology) obeys the AGM postulates. It follows then that the ranking

¹²The generalization to more than two pieces of evidence is straightforward.

system proposed in this paper, also obeys those postulates under the same interpretation of beliefs. Boutilier [3] has further shown that a version¹³ of the priority function Z^+ is associated with the epistemic entrenchment ordering of the AGM theory.

There are several advantages, however, to our system over those proposed by AGM and Spohn, all stemming from the fact that our revision process revolves around rules, not around the beliefs, rankings or expectations that emanate from those rules. By contrast, the logic of AGM is built around “sets of beliefs” or “sets of expectations” and that of Spohn revolves around the ranking function κ . This difference has both practical and philosophical implications. First, since the number of propositions in one’s belief set is astronomical, and so is the number of worlds, it is a computational necessity to base belief revision on rules, whose number is usually manageable. Second, our system extracts both beliefs and rankings from the content of Δ , and requires no outside specification of belief orderings.

Third, in order to be able to undo belief changes, Spohn’s framework must limit the input information to defeasible or imprecise observations (i.e., α -conditioning with $\alpha < \infty$). In our system, we can accommodate both imprecise and precise observations (corresponding to $\alpha = \infty$) using ordinary conditioning. Given a set of precise observations ϕ , the set of beliefs is defined as those propositions σ for which $\kappa(\neg\sigma|\phi) > 0$. Retraction of obsolete observations can be done by simply removing those observations from ϕ .

Fourth, and perhaps most significantly, our system is capable of responding not merely to empirical observations but also to linguistically transmitted information such as conditional sentences (i.e., if-then rules). For example, suppose someone tells us that Mary too tries to avoid John in parties; we simply add this information to our knowledge base in the form of a new rule, $J \rightarrow \neg M$, recompute Z^+ , and are prepared to respond to new observations or hearsay. In Spohn’s system, where revisions are limited to α -conditioning, one cannot properly revise beliefs in response to conditional statements.¹⁴

5 Conclusions

This paper proposes a belief-revision system that reasons tractably and plausibly with linguistic quantification of both observational reports (e.g., “looks like”) and domain rules (e.g., “typically”). We have shown that the system is semi-tractable, namely, tractable for every sublanguage in which propositional satisfiability is polynomial (Horn expressions, network theories, acyclic expressions, etc.). To the best of our knowledge, this is the first system that

¹³The proof in [3] considers the priorities Z^+ resulting from a *flat* set of rules $\{\varphi_i \xrightarrow{\delta_i} \psi_i\}$, namely one in which all δ_i ’s are 0, as in system-Z [15].

¹⁴Gärdenfors [6] attempts to devise postulates for conditional sentences, but finds them incompatible with the Ramsey test (page 156-160).

reasons with approximate probabilities which offers such broad guarantees of tractability.¹⁵ We expect these results to carry over to the theory of possibility as formulated by Dubois and Prade [5], which has similar features to Spohn's system except that beliefs are measured on the real interval $[0, 1]$. In addition we have shown that, without loss of tractability, the system can also accommodate expressions of imprecise observations, thus providing a good model for weighing the impact of evidence and counter-evidence on our beliefs. Finally, we have shown that the enterprise of belief revision, as formulated in the work presented in [2], can find a tractable and natural embodiment in system- Z^+ , unhindered by difficulties that plagued earlier systems.

References

- [1] Ernest W. Adams. *The Logic of Conditionals*. D.Reidel, Dordrecht, Netherlands, 1975.
- [2] C. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [3] Craig Boutilier. Conditional logics for default reasoning and belief revision. PhD dissertation, University of Toronto, 1992. Also Technical Report 91-1, Department of Computer Science, University of British Columbia.
- [4] W. Dowling and J. Gallier. Linear-time algorithms for testing the satisfiability of propositional Horn formulae. *Journal of Logic Programming*, 3:267–284, 1984.
- [5] D. Dubois and H. Prade. *Possibility Theory: An Approach to Computerized Processing of Uncertainty*. Plenum Press, New York, 1988.
- [6] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, 1988.
- [7] Hector A. Geffner. Default reasoning: Causal and conditional theories. PhD dissertation, University of California Los Angeles, Computer Science Department, Los Angeles, 1989. Also Technical Report TR-137 Cognitive Systems Lab., Forthcoming, MIT Press.
- [8] Moisés Goldszmidt, Paul Morris, and Judea Pearl. A maximum entropy approach to nonmonotonic reasoning. In *Proceedings of American Association for Artificial Intelligence Conference*, pages 646–652, Boston, 1990.
- [9] Moisés Goldszmidt and Judea Pearl. System Z^+ : A formalism for reasoning with variable strength defaults. In *Proceedings of American Association for Artificial Intelligence Conference*, Anaheim, CA, 1991.

¹⁵Whereas most tractability results exploit the topological structure of the knowledge base (hypertrees, or partial hypertrees), ours are topology-independent.

- [10] Moisés Goldszmidt and Judea Pearl. Stratified rankings for causal relations. In *Proceedings of the Fourth International Workshop on Nonmonotonic Reasoning* (forthcoming), Vermont, 1992.
- [11] Daniel Lehmann. What does a conditional knowledge base entail? In *Proceedings of Principles of Knowledge Representation and Reasoning*, pages 212–222, Toronto, 1989.
- [12] John McCarthy. Applications of circumscription to formalizing commonsense knowledge. *Artificial Intelligence*, 28:89–116, 1986.
- [13] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.
- [14] Judea Pearl. Jeffrey’s rule, passage of experience and neo-Bayesianism. In R. Parikh, editor, *Defeasible Reasoning and Knowledge Representation*, pages 121–135. Kluwer Publishers, San Mateo, 1990.
- [15] Judea Pearl. System Z: A natural ordering of defaults with tractable applications to default reasoning. In R. Parikh, editor, *Proceedings of Theoretical Aspects of Reasoning about Knowledge*, pages 121–135. Morgan Kaufmann, San Mateo, CA, 1990.
- [16] P. P. Shenoy. On Spohn’s rule for revision of beliefs. *International Journal of Approximate Reasoning*, 5(2):149–181, 1991.
- [17] Yoav Shoham. Nonmonotonic logics: Meaning and utility. In *Proceedings of International Conference on AI (IJCAL87)*, pages 388–393, Milan, Italy, 1987.
- [18] W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W. L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, pages 105–134. Reidel, Dordrecht, Netherlands, 1987.