

On the Revision of Probabilistic Belief States

Craig Boutilier

Department of Computer Science
University of British Columbia
Vancouver, British Columbia
CANADA, V6T 1Z4

Abstract

In this paper we describe two approaches to the revision of probability functions. We assume that a probabilistic state of belief is captured by a counterfactual probability or Popper function, the revision of which determines a new Popper function. We describe methods whereby the original function determines the nature of the revised function. The first is based on a probabilistic extension of Spohn's OCFs, while the second exploits the structure implicit in the Popper function itself. This stands in contrast with previous approaches that associate a unique Popper function with each absolute (classical) probability function. We also describe iterated revision using these models. Finally, we consider the point of view that Popper functions may be abstract representations of certain types of absolute probability functions, but show that our revision methods cannot be naturally interpreted as conditionalization on these functions.

Keywords: belief revision, probability, Popper functions, iterated revision, conditionals, ordinal conditional functions

1 Introduction

Most theories of belief revision take *belief sets*, or sets of (objective) sentences to be the target of the revision process. A revision function maps one belief set into another in response to some new piece of evidence A . The information that determines the exact nature of this mapping can be represented in various ways — for instance, as an entrenchment relation or a set of conditional beliefs — and together with the belief set constitutes an agent’s epistemic state.

It has frequently been suggested that a belief set can be viewed as an abstraction of a more fine-grained, quantitative epistemic state, namely a *probability function* (p-function) that captures degrees of belief. We take the *top* of the p-function P (i.e., those A such that $P(A) = 1$) to be the agent’s belief set, but allow further discriminations among non-beliefs to be held [12, 22]. Of course, the revision problem in this setting requires somewhat more machinery; but once a method for revising a p-function is in hand, revision of the corresponding belief set comes for free: we simply take the revised belief set to be the top of the revised p-function. Probabilistic revision seems to have received somewhat less attention than its qualitative counterpart, in some measure due to the fact that *Popper functions* (or to use Stalnaker’s [29] more suggestive terminology, *counterfactual probability functions*) provide a rather natural and robust representation of an epistemic state and revision method in such a setting. By now the connections between Popper functions and the qualitative representations of revision functions are well-understood [29, 30, 28, 12, 22] and confirm that the “belief as top of a p-function” perspective is tenable.

The problem of iterated revision requires that a revision function produce not only a new belief set or p-function, but a new epistemic state to guide subsequent revisions. On this count many theories are silent and in general, as pointed out by Harper [18], the problem is a difficult one. However, some proposals for iterated revision in the qualitative setting have been put forth. In this article we examine the extent to which similar considerations can be adapted to iterated probabilistic revision. We pay special attention to Spohn’s [28] *ordinal conditional functions (OCFs)* and the author’s [5, 7] proposal for *minimal conditional (MC) revision*. Spohn’s model imposes additional structure on the usual revision functions requiring that the entrenchment and plausibility of sentences be quantified and that the evidence causing revision have an associated strength. The “probabilized” OCFs we introduce further impose such informational demands on the incorporated evidence that the naturalness of the model must be called into question. In addition, we show that under certain natural interpretations, the updates sanctioned by this model cannot be justified by appeal to conditionalization. We then examine a simpler, more impoverished model that allows the direct (iterated) revision of Popper functions and examine some of its properties.

We begin by introducing the AGM theory of revision, a semantic model of revision functions and Spohn’s notion to OCFs in Section 2. In Section 3 we

discuss the connection to Popper functions and present a probabilistic version of OCFs, as well as giving a quasi-infinitesimal interpretation of these models reminiscent of Adam’s [1] ε -semantics. This interpretation suggests that, more generally, one may be able to interpret Popper functions as an appropriate abstraction or summarization of a classical p-function. We then address the question of iterated revision of probability functions in Section 4. We first discuss the probabilistic OCF model and show that the extension of Spohn’s update method to this case is not straightforward. In particular, certain parameters are required to make sense of this model, parameters whose interpretation is unclear. Furthermore, we show that adopting the view that Popper functions correspond to classical p-functions and that revision corresponds to some form of conditionalization does not help fix these parameters. We then examine the probabilistic version of MC-revision and its difficulties. The net result is a somewhat negative conclusion: the revision of probabilistic belief states is not as straightforward, nor as well-understood, as we might have thought.

2 Non-Probabilistic Belief Revision

We assume an agent to have a deductively closed set of beliefs K taken from some underlying language. For concreteness, we will assume this language to be that of classical propositional logic \mathbf{L} , generated by some set of variables \mathbf{P} , and with an associated consequence operation Cn . To keep the technical details to a minimum we assume that \mathbf{P} is finite, giving rise to a finitary language. This will simplify the discussion of probability functions in the next section. We let W denote the finite set of possible worlds (or valuations) suitable for \mathbf{L} . Any world w that satisfies $A \in \mathbf{L}$ (denoted $w \models A$) is dubbed an A -world, the set of which is denoted $\|A\|$. We also use this notation for sets of sentences S , $\|S\|$ denoting those worlds satisfying each element of S . The identically true and false propositions are denoted \top and \perp , respectively.

2.1 The AGM Model

The *expansion* of K by new information A is the belief set $K_A^+ = Cn(K \cup \{A\})$. The *revision* of K by A is denoted K_A^* . The process of revision requires some care, for A may contradict elements of K . Alchourrón, Gärdenfors and Makinson [2] propose a method for logically delimiting the scope of acceptable revisions. To this end, the AGM postulates below, are maintained to hold for any reasonable notion of revision [12].

(R1) K_A^* is a belief set (i.e. deductively closed).

(R2) $A \in K_A^*$.

(R3) $K_A^* \subseteq K_A^+$.

- (R4) If $\neg A \notin K$ then $K_A^+ \subseteq K_A^*$.
- (R5) $K_A^* = Cn(\perp)$ iff $\models \neg A$.
- (R6) If $\models A \equiv B$ then $K_A^* = K_B^*$.
- (R7) $K_{A \wedge B}^* \subseteq (K_A^*)_B^+$.
- (R8) If $\neg B \notin K_A^*$ then $(K_A^*)_B^+ \subseteq K_{A \wedge B}^*$.

It is clear that the epistemic state of an agent cannot consist of a belief set K alone, for K does not contain the information required to determine its revision. At the very least, an epistemic state might be a $\langle K, * \rangle$, where $*$ is some revision function. Less direct but more natural representations of the revision function (at least applied to K) are possible. Among these are *entrenchment relations* [13, 12] and *conditional belief sets* [12, 5]. We briefly describe entrenchment relations below; but the crucial feature is the fact that they capture the revision policies of the agent, the information necessary to revise K to form K_A^* .

Semantically, the epistemic state of an agent can be captured using a *qualitative revision model*.¹ Assuming a fixed set W of possible worlds, a revision model is a transitive, connected ordering relation \leq over W .

Definition 2.1 A *qualitative revision model* (QRM) over W is any relation $\leq \subseteq W \times W$ such that: a) if $w \leq v$ and $v \leq u$ then $w \leq u$; and b) either $w \leq v$ or $v \leq w$ for all $v, w \in W$.

The interpretation of \leq is as follows: $v \leq w$ iff v is as at least as *plausible* a state of affairs as w . Plausibility is a pragmatic measure that reflects the degree to which one would accept w as a possible state of affairs. If v is more plausible than w , loosely speaking v is “more consistent” with the agent’s beliefs than w . Since \leq is a total preorder, W is partitioned into \leq -equivalence classes, or *clusters* of equally plausible worlds. These clusters are themselves totally ordered by \leq . Thus, \leq can be viewed as a qualitative ranking relation, assigning to each world a degree of plausibility. A *K-QRM* is a QRM that captures the epistemic state of an agent with belief set K . In particular, we require that epistemically possible worlds be more plausible than epistemically impossible worlds, and that all epistemically possible worlds are equally plausible. In other words, K -worlds should be exactly those minimal in \leq :

Definition 2.2 A *K-QRM* is any QRM such that $w \leq v$ for all $v \in W$ iff $w \models K$.

For any $A \in \mathbf{L}$ we define

$$\min(\leq, A) = \{w \in W : w \models A, \text{ and } v \models A \text{ implies } w \leq v \text{ for all } v \in W\}$$

¹The presentation is based on the ordering model and logic described in [9, 6]. In our finite setting, this model has no essential differences from Grove’s [14] system of spheres model.

For any satisfiable A we have $\min(\leq, A) \neq \emptyset$.² Intuitively, $\min(\leq, A)$ is the set of most plausible A -worlds. When A is learned, it is this set of worlds that is adopted as the new set of epistemic possibilities. Thus, K_A^* can be defined semantically as

$$\|K_A^*\| = \min(\leq, A)$$

and $B \in K_A^*$ iff $\min(\leq, A) \subseteq \|B\|$. It is easy to verify (see [14, 9]) that the revision function induced by a K -QRM satisfies postulates (R1)–(R8). Furthermore, any function that satisfies the postulates is representable with such a model.³

Within this setting we can define the relative *degree of surprise* associated with sentences as well as the relative *entrenchment* of beliefs. We say A is at *no more surprising* than B iff $\min(\leq, A) \leq \min(\leq, B)$. Intuitively, this reflects the degree to which an agent is willing to accept A as an epistemic possibility. If A is less surprising than B then $A \wedge \neg B \in K_{A \vee B}^*$. If A and B are both believed, we say A is more entrenched than B iff $\neg B$ is less surprising than $\neg A$. This relation holds when an agent is more willing to give up belief in B than A .

2.2 Ordinal Conditional Functions

Spohn [28] introduced a related but somewhat more detailed model of belief revision based on *ordinal conditional functions* or OCFs. Instead of a simple ordering of plausibility over possible worlds, a world is ranked on an ordinal scale according to its degree of plausibility. Spohn recognized that while a qualitative ranking may be appropriate for revising a belief set, this detailed information may be critical when one considers how an entire epistemic state is to be updated.⁴ To simplify the presentation, we assume that plausibility is measured on an integer scale.

Definition 2.3 An *ordinal conditional function* (OCF) over W is a function $\kappa : W \rightarrow \mathbf{N}$ such that $\kappa^{-1}(0) \neq \emptyset$.

The value $\kappa(w)$ indicates the degree of plausibility of situation w , where $\kappa(w) < \kappa(v)$ indicates that w is more plausible than v . We take the worlds w such that $\kappa(w) = 0$ to be those considered epistemically possible by an agent. The restriction $\kappa^{-1}(0) \neq \emptyset$ ensures that the agent’s belief set is consistent. As with QRMs, only certain OCFs are appropriate for an agent with belief set K .

Definition 2.4 A *K-OCF* is any OCF such that

²If we relax the finiteness restriction, this condition can be guaranteed by imposing a certain type of well-foundedness constraint on \leq .

³More accurately, an AGM revision function is representable by a collection of models $\{\leq\}$, with one ordering for each belief set K . Note that if A is unsatisfiable, $\min(\leq, A) = \emptyset$ and $K_A^* = \mathbf{L}$.

⁴We will elaborate on this point in Section 4.

$$\kappa(w) = 0 \text{ iff } w \models K$$

A K -OCF induces a revision function in the obvious fashion. For any $A \in \mathbf{L}$ we define

$$\min(\kappa, A) = \{w \in W : w \models A, \text{ and } v \models A \text{ implies } \kappa(w) \leq \kappa(v) \text{ for all } v \in W\}$$

Defining $\|K_A^*\| = \min(\kappa, A)$ then determines a revision function that satisfies the AGM postulates. Clearly each K -OCF is equivalent to a unique K -QRM under the definition of K_A^* ; but a K -QRM is equivalent to a great number of K -OCFs.⁵ The added expressive power of OCFs will be exploited when we discuss iterated revision and revision of epistemic states in Section 4.

We can extend the plausibility ranking κ to sentences, defining

$$\kappa(A) = \min\{\kappa(w) : w \models A\}$$

The lower $\kappa(A)$ the less surprising A is, with the least degree of surprise (or epistemic possibility) indicated by $\kappa(A) = 0$. Note that either $\kappa(A)$ or $\kappa(\neg A)$ (or both) are zero, and A is believed iff $\kappa(\neg A) > 0$. The entrenchment of a belief is again the dual of surprise: the greater $\kappa(\neg A)$, the more entrenched belief in A is.

3 Probabilistic Revision Models

QRMs and OCFs are possible representations of an agent's epistemic state. Any such model characterizes a unique belief set K as well as a (single-step) revision policy that determines the revised belief set K_A^* . However, belief sets allow only very coarse distinctions in epistemic attitude toward propositions: they can be accepted ($A \in K$), rejected ($\neg A \in K$) or indeterminate ($A, \neg A \notin K$). One might expect an agent to give more or less credence to certain disbelieved possibilities, to be more disposed to one possibility than another without fully accepting or believing the first. Thus, we may suppose that an agent's belief set is determined by a *probability function* (p-function) $P : \mathbf{L} \rightarrow [0, 1]$, satisfying the following conditions:

1. If $\vdash A \equiv B$ then $P(A) = P(B)$
2. If $\vdash A \supset \neg B$ then $P(A \vee B) = P(A) + P(B)$
3. If $\vdash A$ then $P(A) = 1$

Accepted (rejected) propositions are those A such that $P(A) = 1$ ($P(A) = 0$). However, indeterminate propositions are now graded according to their

⁵This distinction is analogous to that made by Spohn [28] between OCFs and SCFs.

probability.⁶ The function $P^\perp(A) = 1$ for all $A \in \mathbf{L}$ is dubbed the *inconsistent p-function* and will sometimes be treated as a p-function (corresponding to the belief set $K^\perp = \mathbf{L}$).

Rather than taking belief sets as primitive, we assume that an epistemic state contains a p-function from which a belief set K is derived.

Definition 3.1 A p-function P is *compatible* with a belief set K just when $P(A) = 1$ iff $A \in K$.

Each p-function is compatible with a unique (deductively closed) belief set K . Since many p-functions are compatible with a fixed K , we take p-functions to be the basic notion from which belief sets are derived. This corresponds to the well-used tactic (e.g., [12, 22]) of defining a belief set to be the *top* of a p-function. We define the conditional probability $P(B|A)$ as $\frac{P(A \wedge B)}{P(A)}$ if $P(A) \neq 0$; by convention we set $P(B|A) = 1$ for all B otherwise.

Semantically, a p-function can be characterized by a (normalized) *weighting function* $P : W \rightarrow [0, 1]$ such that $\sum\{P(w) : w \in W\} = 1$. This induces a p-function (over \mathbf{L}) via the standard relationship:

$$P(A) = \sum\{P(w) : w \models A\}$$

We use P to denote both the weighting function and the induced p-function since each uniquely determines the other. We will also use unnormalized weighting functions, which assign arbitrary positive weights to worlds. An unnormalized function P' determines a normalized function P as follows:

$$P(w) = \frac{P'(w)}{\sum\{P'(v) : v \in W\}}$$

We note that if P is compatible with K then $P(w) > 0$ iff $w \models K$. So it is precisely the epistemically possible worlds that are accorded positive probability.

3.1 Counterfactual Probability Models

The notion of compatibility can be extended to QRMs in the obvious way: we say a p-function P is compatible with $\text{QRM } \leq$ iff P is compatible with the belief set K induced by \leq . Since the epistemic possibilities given by K correspond to the set of worlds minimal in \leq , an appropriate p-function can be imposed by a weighting function P with the property that $P(w) > 0$ iff $w \in \min(\leq, \top)$. An agent's epistemic state might then be taken to consist of a $\text{QRM } \leq$ together with a compatible p-function P . But, while \leq is sufficient to determine the content of K_A^* , a revised epistemic state must include a revised p-function P_A^* ;

⁶Note that notions such as entrenchment or plausibility of sentences do not make any such distinction between indeterminate propositions.

and \leq does not contain the required information. Some method for revising p-functions is needed.

In the special case where $\neg A \notin K$ (or equivalently, $P(A) > 0$), we can use conditionalization of P by A to effect revision by A and derive a revised p-function. We simply set $P_A^* = P(\cdot|A)$. The following observation is straightforward (see, e.g., [12]).

Proposition 3.1 *Let \leq be a K-QRM such that $\neg A \notin K$, let P be a p-function compatible with \leq , and let $P_A^* = P(\cdot|A)$. Then P_A^* is compatible with K_A^* .*

Thus “consistent” revision and conditionalization correspond in the desired manner. If $\neg A \in K$ then $P(A) = 0$ and defining P_A^* via conditionalization results in the inconsistent p-function P^\perp . To alleviate this difficulty, we introduce nonstandard conditional p-functions, or *Popper functions* [29, 30, 12, 22]. A Popper function is a mapping $P : \mathbf{L} \times \mathbf{L} \rightarrow [0, 1]$ satisfying [22]:

1. $P(A \uparrow A) = 1$
2. $P(\cdot \uparrow A)$ is a p-function
3. If $\vdash A \equiv B$ then $P(C \uparrow A) = P(C \uparrow B)$
4. $P(A \wedge B \uparrow C) = P(A \uparrow C) \cdot P(B \uparrow A \wedge C)$

An absolute p-function is defined by setting $P(A) = P(A \uparrow \top)$.⁷

By taking conditional probability as the primitive relation we can impose nontrivial constraints on the value of $P(B \uparrow A)$ even when $P(A) = 0$, and revision of a p-function can be defined by taking $P_A^*(B) = P(B \uparrow A)$. The relationship between the revision of p-functions using nonstandard conditionalization and the AGM revision of belief sets is quite close (see, e.g., [12, Ch.5] or [22]). We take revision of a p-function P as defined above to be basic, assuming P is determined by an appropriate Popper function. This induces an revision function on belief sets, where K is the top of P and K_A^* is the top of P_A^* . Such qualitative revision functions satisfy the AGM postulates, suggesting that the Popper functions are an appropriate representation of probabilistic revision functions.

This relationship can be understood from a semantic perspective as well. The semantics of simple p-functions uses probability weights on epistemically possible worlds (consistent with the induced belief set K). This can be extended to Popper functions by associating weights with *all* worlds W , regardless of their plausibility ranking according to \leq . In this way, the relative probability of worlds within the set $\min(\leq, A)$ is specified and a compatible p-function P_A^* can be derived. The weights of worlds in the same cluster of \leq captures their relative likelihood *should the agent accept them as epistemically possible*. They can be viewed as *counterfactual probabilities* in the sense of Stalnaker [29].

⁷We will consistently use $P(\cdot|\cdot)$ to denote standard conditional probability functions and $P(\cdot \uparrow \cdot)$ to denote nonstandard, Popper functions.

Definition 3.2 $\leq^P = \langle \leq, P \rangle$ is a *counterfactual probability model* (CPM) iff

- (a) \leq is a QRM; and
- (b) $P : W \rightarrow (0, 1]$

Definition 3.3 Let $\leq^P = \langle \leq, P \rangle$ be a CPM. The *counterfactual probability* of B given A (relative to \leq^P) is

$$P(B \uparrow A) = \frac{\sum \{P(w) : w \in \min(\leq, A) \text{ and } w \models B\}}{\sum \{P(w) : w \in \min(\leq, A)\}}$$

Definition 3.4 The *factual probability* of A (relative to M) is

$$P(A) = P(A \uparrow \top)$$

Clearly, the factual probability function P is a p-function. We take the unconditional, factual probability function P to define the objective epistemic state of the agent in the usual way. This factual p-function is compatible with the QRM component of the CPM.

Proposition 3.2 Let $\leq^P = \langle \leq, P \rangle$ be a CPM such that \leq determines belief set K . Then $P(A) = 1$ iff $A \in K$.

We define factual conditional probability in the usual way for P , and denote this with the usual conditioning bar:

Definition 3.5 $P(B|A) = \frac{P(A \wedge B)}{P(A)}$ for all A such that $P(A) > 0$.

We can now describe the new p-function P_A^* that results when the agent's original epistemic state P is revised by A . This revision will proceed by means of *counterfactual conditionalization*.

Definition 3.6 Let P be the factual p-function determined by M . The *revised factual probability function* P_A^* is given by

$$P_A^*(B) = P(B \uparrow A)$$

The (factual) p-function of an agent after such a revision is P_A^* . The following results are easy to verify, and are simple restatements of well-known facts. We emphasize them as they indicate that the process of counterfactual conditionalization conforms to the rationality constraints imposed by our original qualitative considerations.

Proposition 3.3 If A is satisfiable then P_A^* is a (consistent) p-function.

Theorem 3.4 The p-function P_A^* is compatible with the belief set K_A^* .

Theorem 3.5 *If $P(A) > 0$ then $P_A^*(B) = P(B|A)$.*

Thus CPMs are consistent extensions of the AGM theory to the probabilistic case, just like Popper functions. That CPMs in fact determine Popper functions can be seen by appeal to the representation result of van Fraassen [30]. He demonstrates that Popper functions can be represented by an *ordinal family of p-functions*, or a sequence of p-functions ranked according to plausibility. In our finitary setting, we can use a finite ordered family $\{P_1, \dots, P_n\}$ of p-functions over \mathbf{L} . The *minimal A-permitting p-function* P^A is the first P_i in the sequence that accords A positive probability; that is, $P^A = P_i$ where $P_i(A) > 0$ and $P_j(A) > 0$ only if $j \geq i$. We can use this family to define a counterfactual probability function as follows: if $P_i(A) > 0$ for some i , we define $P(B\uparrow A) = P^A(B|A)$; if $P_i(A) = 0$ for all i , we call A an *abnormal* proposition and set $P(B\uparrow A) = 1$ for all B . van Fraassen shows that any such conditional operator is a Popper function and that any Popper function is representable by such a family.

Indeed, the stronger representation result of Spohn [27] can be adopted in our finitary setting. Following Spohn, we say an ordered family is *dimensional* iff for each P_i there is a sentence A_i such that $P_i(A_i) = 1$ and $P_j(A_i) = 0$ for all $j < i$. If all p-functions are σ -additive, dimensional ordered families can be used to represent Popper functions. Dimensional ordered families have several nice properties, including *minimality* in the sense that the ordered family of p-functions cannot be replaced by a smaller family. Indeed, a dimensional family satisfies an even stronger minimality requirement, for the elements of P_i are *orthogonal*: if $i \neq j$ then $P_i(A) = 1$ and $P_j(A) = 0$ for some A [12]. Thus, not only is the family itself as small as possible, its elements are as well.

A simple reconstruction of CPMs demonstrates that they are equivalent to finite ordered families of p-functions and thus equivalent to Popper functions. While rather straightforward, we spell out the connection in detail since we will exploit the correspondence frequently below. For any QRM \leq , let $W_1 \cdots W_n$ denote the equivalence classes of W determined by \leq . That is, for each W_i , if $w, v \in W_i$ then $w \leq v$; and if $w \in W_i$, $v \leq w$ and $w \leq v$ then $v \in W_i$. Furthermore, assume that if $i < j$ then there is some $w \in W_i, v \in W_j$ such that $w < v$. Let \leq^P be a CPM. For each equivalence class W_i define a (normalized) weighting function (and equivalent p-function) P_i as

$$P_i(w) = \frac{P(w)}{\sum\{P(v) : v \in W_i\}}$$

if $w \in W_i$ and $P_i(w) = 0$ if $w \notin W_i$. We call $\{P_i : i \leq n\}$ the ordered family of p-functions induced by \leq^P . A counterfactual probability function can be defined using the relationship:

$$P(B\uparrow A) = P^A(B|A)$$

It is easy to verify that this counterfactual probability function is exactly that determined by the original CPM. Furthermore, given any such ordered family of p-functions, it is easy to construct a corresponding CPM. The results of van Fraassen and Spohn ensure that CPMs determine Popper functions and that any Popper function is representable by a CPM.⁸ Furthermore, the ordered family generated is dimensional — it is therefore minimal and consists of pairwise orthogonal elements.

In the sequel we will use the original definition of CPMs and their representation as a minimal, orthogonal ordered family of p-functions interchangeably.

3.2 Probabilistic OCFs

Just as we probabilized QRMs by adding probability weights to the worlds in the qualitative ranking, we can probabilize OCFs.⁹

Definition 3.7 $\kappa_P = \langle \kappa, P \rangle$ is a *probabilistic ordinal conditional function* (POCF) iff

- (a) κ is an OCF; and
- (b) $P : W \rightarrow (0, 1]$

The Popper function $P(\cdot|\cdot)$ induced by a POCF is defined in exactly the same fashion as for CPMs. A POCF determines a minimal, orthogonal ordered family of p-functions in precisely the same way as CPMs. However, we will index the elements of this ordered family by the κ -ranking associated with the worlds over which it is defined; that is, the ordered family will be written

$$\{P_i : \kappa^{-1}(i) \neq \emptyset\}$$

While POCFs and Popper functions correspond in the obvious way, we see that many different POCFs are equivalent to the same CPM and induce the same Popper function. Thus we might think of POCFs as a Popper function with additional structure.

3.3 A Standard Interpretation of POCFs

If one is going to use probabilities as degrees of belief, it seems natural to question the need for Popper functions, ordered families of p-functions and (categorical) belief sets. If one is going to allow a sentence A in K to be retracted

⁸We should point out that since CPMs assign non-zero weight to every element of W , they can only represent Popper functions such that no consistent A is abnormal; that is, we must have $P(\neg A|A) = 0$. To capture all Popper functions, we can simply relax the requirement of non-zero weight, thus allowing abnormal elements. We will not be concerned with such functions in the sequel.

⁹Indeed, in his original paper Spohn suggests that OCFs could be probabilized.

when $\neg A$ is learned, why not simply assign A some degree of belief less than 1 in the first place and use standard techniques such as conditionalization to incorporate new items of belief?

If one wishes to allow the possibility that any “belief” can be overturned given the proper evidence, then full belief can be granted only to tautologies, and every contingency must have some positive probability. To take a slightly less extreme view, one might accord observational reports (say) the status of full belief, but still no conclusions drawn from these would be certain. Presumably, there are certain computational advantages to be gained by ruling out possibilities that are very unlikely [10, 17]. Chief among these is the ability to exploit logical rules of inference. Such rules allow conclusions to be reached in manner that is independent of context, in contrast to probabilistic inference. The *locality* of logical rules can be exploited if parts of the belief are (treated as if they are) fully belief [23].

It may also be that the cost associated with reaching incorrect (unhedged) conclusions and being forced to revise the belief set is outweighed by the probability of being correct. We might therefore think of a constraint $B \in K_A^*$ as an instantiation of an *acceptance rule* [20]. If $B \in K_A^*$ is satisfied by the epistemic state of an agent, we take it to mean that there is a certain utility associated with complete acceptance of B given A .¹⁰ On this view, it is reasonable to allow a conditional belief $B \in K_A^*$ to be held even when $\neg A$ is accorded full belief, $P(\neg A) = 1$. Consequently, we do not take a $P(A) = 0$ to indicate that A is (logically or physically) impossible, but simply that it is not, to use Levi’s [21] terminology, a *serious possibility*.

To make sense of this perspective, it should be possible to interpret a Popper function, a CPM or a POCF as some sort of abstraction of a classical, absolute p-function. There should be some “true” p-function \overline{P} such that the Popper function P induces the appropriate beliefs with respect to \overline{P} . Furthermore, if P is representable by some minimal, orthogonal ordered family $\{P_i\}$ of p-functions, the true p-function \overline{P} should be constructed through some combination of the elements P_i . In particular, we expect \overline{P} to be some *additive mixture* of the P_i ; that is,

$$\overline{P}(A) = a_1 \cdot P_1(A) + \dots + a_n \cdot P_n(A)$$

Of course, not any additive mixture will do. We expect a mixture to justify in some way the “acceptance rules” implicit in the Popper function. Below we suggest one such interpretation, whereby if some proposition is more plausible than another, the first can be made arbitrarily more probable than the second with respect to \overline{P} . We remark that other interpretations are possible that can be modeled using an additive mixture of the family $\{P_i\}$.¹¹ We do not suggest

¹⁰We do not address here the issue of how one determines appropriate acceptance rules. Below we show how such acceptance rules can be interpreted in a very strong probabilistic way that guarantees this to be the case; but in general decision-theoretic criteria should be brought to bear [24, 8].

¹¹One example is using standard acceptance rules so that $P_A^*(B) = k$ ensures that $P(B|A) >$

that this interpretation is the correct way to view Popper functions, but simply present it as an alternative to illustrate the feasibility of this point of view.

Imagine an agent whose epistemic state is represented as a POCF which, by our standard construction, determines an ordered family $\{P_i : \kappa^{-1}(i) \neq \emptyset\}$. We suppose that P is an absolute p-function abstracted by this POCF. The serious possibilities admitted by the agent are those sentences A such that $\kappa(A) = 0$. If A is a serious possibility and B is not, we should expect that A is more probable than B by some significant factor, for instance, $P(A) \cdot \varepsilon \geq P(B)$ for some sufficiently small $\varepsilon > 0$. Furthermore, we should expect that the degree of plausibility of a proposition (its κ -rank) determines the extent of this difference. In general, we require that if $\kappa(A) + i = \kappa(B)$ for some $i > 0$ then $P(A) \cdot \varepsilon^i \geq P(B)$. That is, more plausible sentences can be made arbitrarily more probable than less plausible sentences, and difference in plausibility forces a “lower bound” on this difference. For any POCF, such a p-function P can be constructed as an additive mixture of its ordered family.

Theorem 3.6 *Let $\langle \kappa, P \rangle$ be a POCF determining an ordered family $\{P_i : \kappa^{-1}(i) \neq \emptyset\}$. For any sufficiently small ε , there exists a p-function $\overline{P} = \sum \{a_i \cdot P_i\}$ such that if $\kappa(A) + i = \kappa(B)$ for some $i > 0$ then $\overline{P}(A) \cdot \varepsilon^i \geq \overline{P}(B)$.*

Proof For simplicity, we assume that ε satisfies the rather weak constraint that $\varepsilon^i \geq \sum_{j>i} \varepsilon^j$ (although for somewhat larger ε the construction can be modified). To prove the result we must determine appropriate parameters a_i . Let $W_i = \{w : \kappa(w) = i\}$. We have by construction that $P_i(w) > 0$ iff $w \in W_i$, where P_i denotes the (normalized) weighting function corresponding to p-function P_i . For each P_i define

$$\min(P_i) = \min\{P_i(w) : w \in W_i\}$$

Note that $\min(P_i) \leq 1$. We define the (unnormalized) additive parameters as follows:

$$a_j = \varepsilon^{2j} \cdot \prod_{i<j} \{\min(P_i)\}$$

The p-function \overline{P} is represented by the weighting function (where $w \in W_j$):

$$\overline{P}(w) = P_j(w) \cdot \varepsilon^{2j} \cdot \prod_{i<j} \{\min(P_i)\}$$

Now suppose $\kappa(A) = k$ and $\kappa(B) = k + i$. In the following, P is unnormalized since the additive parameters a_i are unnormalized; but all relationships hold when normalization is performed. Recall that $\min(P_i) \leq 1$.

$k \cdot (1 - \varepsilon)$ for some small $\varepsilon > 0$. A variant of Adam’s [1] semantics for conditionals can be used to verify that such an interpretation can be given and that it corresponds to an additive mixture (where the weights a_i are a function of ε).

Since $\kappa(A) = k$, $W_k \cap \|A\| \neq \emptyset$. So

$$\begin{aligned}\overline{P}(A) &\geq \min(P_k) \cdot \varepsilon^{2k} \cdot \prod_{j < k} \{\min(P_j)\} \\ &\geq \min(P_k) \cdot \varepsilon^{2k} \cdot \prod_{j \leq k} \{\min(P_j)\}\end{aligned}$$

Since $\kappa(B) = k + i$, $W_j \cap \|B\| = \emptyset$ for all $j < k + i$. So

$$\begin{aligned}\overline{P}(B) &\leq \sum_{j=k+i}^n \{\varepsilon^{2j} \cdot \prod_{l < j} \{\min(P_l)\}\} \\ &\leq \prod_{j \leq k} \{\min(P_j)\} \cdot \sum_{j=k+i}^n \{\varepsilon^{2j} \cdot \prod_{l=k+1}^{j-1} \{\min(P_l)\}\} \\ &\leq \prod_{j \leq k} \{\min(P_j)\} \cdot \sum_{j=k+i}^n \{\varepsilon^{2j}\} \\ &\leq \prod_{j \leq k} \{\min(P_j)\} \cdot \varepsilon^{2k+i} \\ &\leq \prod_{j \leq k} \{\min(P_j)\} \cdot \varepsilon^{2k} \cdot \varepsilon^i \\ &\leq \overline{P}(A) \cdot \varepsilon^i\end{aligned}$$

■

Thus for any small $\varepsilon > 0$, a suitable p-function can be constructed that validates the relationship imposed by the POCF. The p-function constructed in this proof also has the property that if $\kappa(A) = i$ then $P(A) \leq \varepsilon^i$.

Proof If $\kappa(A) = i$ then

$$\overline{P}(A) \leq \frac{\prod_{j=i}^n \{\varepsilon^{2j}\}}{\prod_{j=0}^n \{\varepsilon^{2j}\}}$$

Thus,

$$\frac{\overline{P}(A)}{\varepsilon^i} \leq \frac{\prod_{j=i}^n \{\varepsilon^{2j}\}}{\prod_{j=0}^n \{\varepsilon^{i+2j}\}}$$

So $\frac{\overline{P}(A)}{\varepsilon^i} \leq 1$ and $\overline{P}(A) \leq \varepsilon^i$. ■

Thus a POCF can be viewed as an abstract representation of some reasonable additive mixture of its corresponding ordered family of p-functions. Again, other reasonable mixtures are possible — the crucial point is that classical interpretations of POCFs based on additive mixtures of their ordered family presentations are feasible.

4 Iterated Revision

One difficulty faced by the revision models presented to this point is that they do not lend themselves to iteration. In the nonprobabilistic case an epistemic state is represented by a QRM or OCF suitable for some belief set K , and describes the form to be taken by a new belief set K_A^* . But nothing in these models suggests what form should be taken by the new epistemic state, a new QRM or OCF suitable for K_A^* . This new model is necessary to process subsequent revisions of K_A^* . Similar considerations apply to probabilistic models, be they CPMs (simple Popper functions) or PO CFs (structured Popper functions) suitable for some initial p-function P . While they dictate the precise form of P_A^* , they provide no guidance for the construction of a new probabilistic epistemic state (a CPM or PO CF). This difficulty was first noticed by Harper [18].

In the nonprobabilistic setting, Gärdenfors [12] circumvented this difficulty by introducing the notion of a *belief revision system*, a set \mathbf{K} of belief sets and an AGM revision function $*$ that maps $\langle K, A \rangle$ (where $K \in \mathbf{K}$ and $A \in \mathbf{L}$) to $K_A^* \in \mathbf{K}$. Since $*$ is defined for all $K \in \mathbf{K}$, the revision of K_A^* is determined and iterated revision proceeds unobstructed. Such a model of iteration is unattractive for two reasons. First, the revision of a belief set K is fixed. An agent with objective belief set K cannot accept different revision policies with respect to K (for example, at different times). Since there are many QRMs and OCFs suitable for any fixed K , there seems little reason to expect an agent to choose one such epistemic state once and for all. The second undesirable feature of this model is its generality. The revision of K is dictated by a plausibility ordering \leq suitable for K . A belief revision system requires that the revision of K_A^* also be modeled by (or representable as) a new plausibility ordering \leq_A^* .¹² However, it imposes no constraints on the relation between \leq and \leq_A^* other than what we dub the *Basic Requirement*.

The Basic Requirement: If \leq is a K -QRM determining revision function $*$, then the revision model \leq_A^* must be such that $\min(\leq_A^*, \top) = \min(\leq, A)$.

Thus, essentially arbitrary changes in the plausibility ordering \leq are permitted. Equivalently, an agent is permitted to make arbitrary changes in its judgements of the relative entrenchment and plausibility of propositions. A model of iterated revision that imposes additional structure on the change of epistemic state is therefore desirable.

Lindström and Rabinowicz [22] have proposed a similar model for the revision of p-functions. They suggest that each p-function be associated with a unique Popper function. The Popper function associated with p-function P determines the revised p-function P_A^* as described in the last section. Subsequent revision of P_A^* is determined by its corresponding Popper function. The unattractive features of belief revision systems are inherited by this probabilistic

¹²Analogous remarks hold for the revision of OCFs.

model of iteration: each p-function is bound to a single revision policy (Popper function) and the relationship between a Popper function and its revision can be arbitrary.

4.1 Iteration using POCFs

As we hinted in Section 2.3, the integer degrees of plausibility associated with worlds and propositions in Spohn's OCFs play a crucial role in iterated revision. Indeed, Spohn's model of revision, in contrast with the AGM theory, is influenced primarily by the problem of iteration, with revision by A defined so that a new OCF is produced. Our presentation of the revised belief set K_A^* determined by an OCF κ is a mere by-product of Spohn's model of belief dynamics, which we now present.

In OCFs beliefs are held with a specified degree of strength. If $\neg A \in K$, then the strength of belief in $\neg A$ is $\kappa(A)$. If revision of K by A is not only to produce a new belief set K_A^* , but also a new OCF $\kappa_{A,k}^*$, then the strength with which A is to be held in the revised OCF must be given, for otherwise the revision will be underdetermined. In a manner reminiscent of Jeffrey conditionalization [19], we define the A, k -conditionalization of κ as follows.

Definition 4.1 Let κ be an OCF, $A \in \mathbf{L}$ be satisfiable, and $k \geq 0$. The A, k -conditionalization of κ is the OCF $\kappa_{A,k}$ where

$$\kappa_{A,k}(w) = \begin{cases} \kappa(w) - \kappa(A) & \text{if } w \in \|A\| \\ k + (\kappa(w) - \kappa(\neg A)) & \text{if } w \in \|\neg A\| \end{cases}$$

If κ is an OCF such that $\neg A$ is believed, intuitively κ is adjusted by shifting the A -worlds down by a factor of $\kappa(A)$ — this ensures that $\min(\kappa, A)$ is given a new ranking of 0 (down from $\kappa(A)$) — and by shifting all $\neg A$ -worlds up by a factor of k (from 0) — ensuring that A is now believed to degree k . Since worlds within the sets $\|A\|$ and $\|\neg A\|$ stand in the same relation, each new cluster within $\kappa_{A,k}$ will typically be the union of the some set of A -worlds of some fixed rank (w.r.t. κ) and some set of $\neg A$ -worlds of a different fixed rank. In particular, we have:

$$\kappa_{A,k}^{-1}(i) = \{w \in \|A\| : \kappa(w) = i + \kappa(A)\} \cup \{w \in \|\neg A\| : \kappa(w) = i + k - \kappa(\neg A)\}$$

Both sets may be empty in which case no worlds are assigned rank i in the revised OCF. If not, we classify the resulting cluster as follows: if $\kappa_{A,k}^{-1}(i) \subseteq \|A\|$ we say that $\kappa_{A,k}^{-1}(i)$ is a *resulting A-cluster*; if $\kappa_{A,k}^{-1}(i) \subseteq \|\neg A\|$ it is a *resulting $\neg A$ -cluster*; otherwise, it is a *mixed cluster*. Mixed clusters will play an important role in the revision of POCFs.

It is easily seen that if $k > 0$, this is a revision operation that accepts A .¹³ Note that to effect revision by A in a way that defines a new OCF $\kappa_{A,k}(w)$ a

¹³The restriction that $k \geq 0$ is made for convenience of presentation. The A, k -conditionalization of κ can be captured by $\neg A, -k$ -conditionalization if $k < 0$.

strength of evidence k must be associated with the new fact A . Clearly, the revision process can be iterated, for the A, k -conditionalization of κ gives a new OCF suitable for the revised belief set K_A^* . In other words, A, k -conditionalization satisfies the Basic Requirement that $\min(\kappa_{A,k}, \top) = \min(\kappa, A)$. Furthermore, the new OCF is strongly related to the original OCF. An agent's new judgements of plausibility, surprise and entrenchment correspond in a natural way to its original epistemic state.

We now examine the extent to which Spohn's A, k -conditionalization operation can be applied to POCFs and provide us with a model of probabilistic revision which deals with iteration. We first point out that a probabilistic extension of A, k -conditionalization cannot be viewed as a means of revising standard Popper functions. As we noted earlier, POCFs have additional structure and this structure is exploited by A, k -conditionalization. Thus, the view of Popper functions as an appropriate representation of a probabilistic epistemic state in untenable if we use this revision model to effect changes in epistemic state. Spohn [28], following Harper, indeed suggests that Popper functions are impoverished for precisely this reason.

Let κ^P be a POCF with underlying OCF κ and weighting function P . The Popper function $P(\cdot \uparrow \cdot)$ induced by κ^P is compatible with κ in the sense that its ordered family presentation $\{P_i : \kappa^{-1}(i) \neq \emptyset\}$ directly corresponds to the clusters of κ : we have $P_i(w) > 0$ iff $\kappa(w) = i$. When we perform A, k -conditionalization, we would like the revised POCF $\kappa_{A,k}^P$ and its ordered family representation to stand in the same correspondence with the revised OCF $\kappa_{A,k}$.

Definition 4.2 Let $\kappa^P = \langle \kappa, P \rangle$ be a POCF. A POCF $\kappa_{A,k}^P$ is an A, k -conditionalization of κ^P iff $\kappa_{A,k}^P = \langle \kappa_{A,k}, P' \rangle$ where P' is an arbitrary (positive) weighting function.

If $P(\cdot \uparrow \cdot)$ is the Popper function induced by κ^P , we use $P^{A,k}(\cdot \uparrow \cdot)$ to denote the Popper function corresponding to some A, k -conditionalization $\kappa_{A,k}^P$. We also use $P_i^{A,k}$ to denote elements of the ordered family of p-functions induced by $\kappa_{A,k}^P$.

Such arbitrary changes leave $\kappa_{A,k}^P$ drastically underspecified. More importantly, this class of change functions admits some unintuitive changes in epistemic state. At the very least, when $k > 0$ we should require that $P_0^{A,k} = P(\cdot \uparrow A)$ — revising by A should induce the same objective state of belief as counterfactual conditionalization on A . Furthermore, in analogy with Jeffrey conditionalization, we might require that an update by A not change the conditional probabilities within the A -part or the $\neg A$ -part of P ; in other words, we insist that $P^{A,k}(B \uparrow A) = P(B \uparrow A)$ and $P^{A,k}(B \uparrow \neg A) = P(B \uparrow \neg A)$ for all B . Neither of these restrictions is enforced by arbitrary A, k -conditionalization of κ^P . These conditions can be captured by insisting that the new weighting component of $\kappa_{A,k}^P$ keep the relative weights of worlds within the A -part (respectively, the $\neg A$ -part) of each new cluster fixed. For resulting A -clusters (resp. $\neg A$ -clusters)

the weighting function restricted to that cluster corresponds to the conditional probability function $P_i(\cdot|A)$ (resp. $P_i(\cdot|\neg A)$) for some i . For resulting mixed clusters, the p-functions for both the A and $\neg A$ -parts must be mixed in an appropriate way. We define *regular A, k -conditionalizations* to capture this notion.

Definition 4.3 Let P_i and P_j be two p-functions. The p-function P is a non-trivial δ -mixture of $\langle P_i, P_j \rangle$ iff $P(A) = \delta P_i(A) + (1 - \delta)P_j(A)$, where $0 < \delta < 1$.

Definition 4.4 Let $\kappa^P = \langle \kappa, P \rangle$ be a POCF and $m = \max\{i : \kappa_{A,k}^{-1}(i) \neq \emptyset\}$. Let $\delta_0, \dots, \delta_m$ be a set of *mixture factors*, $0 < \delta_i < 1$. We say $\kappa_{A,k}^P = \langle \kappa_{A,k}, P^{A,k} \rangle$ is a *regular A, k -conditionalization with mixture factors* $\delta_0, \dots, \delta_m$ of κ^P iff $\kappa_{A,k}^P$ is an A, k -conditionalization of κ^P and $P^{A,k}$ satisfies the following properties:

1. if $\kappa_{A,k}^{-1}(i)$ is a mixed cluster and $\kappa_{A,k}(w) = i$ then:
 - (a) if $w \in \|A\|$ then $P^{A,k}(w) = \frac{\delta_i P(w)}{\sum\{P(v) : \kappa_{A,k}(v) = i, v \in \|A\|\}}$; and
 - (b) if $w \in \|\neg A\|$ then $P^{A,k}(w) = \frac{(1 - \delta_i)P(w)}{\sum\{P(v) : \kappa_{A,k}(v) = i, v \in \|\neg A\|\}}$
2. if $\kappa_{A,k}^{-1}(i)$ is not a mixed cluster and $\kappa_{A,k}(w) = i$ then $P^{A,k}(w) = P(w)$

Intuitively, a regular A, k -conditionalization of κ^P provides a well-behaved weighting function over $\kappa_{A,k}$. If a cluster induced by $\kappa_{A,k}$ is not mixed (that is, it consists solely of A -worlds or $\neg A$ -worlds), the relative weights of worlds within that cluster are unchanged from κ^P . For a mixed cluster $\kappa_{A,k}^{-1}(i)$, the A -worlds retain the same relative weight but the total weight is fixed by the some mixture factor δ_i ; and the total weight of $\neg A$ -worlds is fixed by its complement $1 - \delta_i$. We have assumed the existence of the factors for each cluster in $\kappa_{A,k}$ (as well as “empty clusters”), but only the δ_i corresponding to mixed clusters are used and need be specified.

For a fixed set of mixture factors it is clear that the regular A, k -conditionalization of κ^P is unique and induces a fixed Popper function. This Popper function satisfies the properties we expect.

Theorem 4.1 Let κ^P be a POCF and $P(\cdot \uparrow \cdot)$ its corresponding Popper function. Let $\kappa_{A,k}^P$ be a regular A, k -conditionalization of κ^P and $P^{A,k}(\cdot \uparrow \cdot)$ its corresponding Popper function. Then $P(B \uparrow A \wedge C) = P^{A,k}(B \uparrow A \wedge C)$ and $P(B \uparrow \neg A \wedge C) = P^{A,k}(B \uparrow \neg A \wedge C)$ for all B, C .

Proof Assume $A \wedge C$ is satisfiable (otherwise $P(B \uparrow A \wedge C) = P^{A,k}(B \uparrow A \wedge C) = 1$ for all B). We have

$$P(B \uparrow A \wedge C) = \frac{\sum\{P(w) : w \in \min(\kappa, A \wedge C) \text{ and } w \models B\}}{\sum\{P(w) : w \in \min(\kappa, A \wedge C)\}}$$

We also have $\min(\kappa_{A,k}, A \wedge C) = \min(\kappa, A \wedge C)$. If $\min(\kappa_{A,k}, A \wedge C)$ is in a resulting A -cluster then $P^{A,k}(w) = P(w)$ for all $w \in \min(\kappa_{A,k}, A \wedge C)$. Otherwise, it is in a mixed cluster and $P^{A,k}(w) = \delta_i P(w)$. Thus, the relative weights of all worlds in $\min(\kappa_{A,k}, A \wedge C)$ are unchanged and $P(B \uparrow A \wedge C) = P^{A,k}(B \uparrow A \wedge C)$. An analogous argument can be made for $P^{A,k}(B \uparrow \neg A \wedge C)$. ■

Corollary 4.2 $P(B \uparrow A) = P^{A,k}(B \uparrow A)$ and $P(B \uparrow \neg A) = P^{A,k}(B \uparrow \neg A)$.

Corollary 4.3 If $k > 0$ then $P^{A,k}(B \uparrow \top) = P(B \uparrow A)$.

Gärdenfors [12] has proposed a set of postulates for the revision of p-functions that mirror the postulates (R1)–(R8). We present a modified version of these.¹⁴ Let P be a p-function and A some consistent sentence. A probability revision operation satisfies:

- (P1) P_A^* is a (consistent) p-function.
- (P2) $P_A^*(A) = 1$.
- (P3) If $P(A) > 0$ then $P_A^* = P(\cdot|A)$
- (P4) If $\models A \equiv B$ then $P_A^* = P_B^*$.
- (P5) If $P_A^*(B) > 0$ then $P_{A \wedge B}^* = P_A^*(\cdot|B)$.

Given the previous results, it is quite easy to verify that regular A, k -conditionalization of κ^P determines a probability revision operation satisfying postulates (P1)–(P5) if we take P to be the absolute p-function induced by κ^P and P_A^* to be the p-function induced by $\kappa_{A,k}^P$ (for any $k > 0$).

Regular A, k -conditionalization of κ^P is defined using the POCF explicitly and constructing an unnormalized weighting function. However, the same effect can be achieved by applying a similar operation to the ordered family presentation of κ^P to form a revised ordered family. In some sense this operation may be more natural as operates on the canonical representation of POCFs.

Definition 4.5 Let $\{P_i : \kappa^{-1}(i) \neq \emptyset\}$ be the ordered family of p-functions induced by some POCF κ^P and let $\delta_0, \dots, \delta_m$ be a set of mixture factors, $0 < \delta_i < 1$. The regular A, k -conditionalized family of p-functions $\{P_i^{A,k} : \kappa_{A,k}^{-1}(i) \neq \emptyset\}$ is defined as follows:

$$P_i^{A,k} = \begin{cases} \delta_i P_{\kappa(A)+i}(\cdot|A) + (1 - \delta_i) P_{i+k-\kappa(\neg A)}(\cdot|\neg A) & \text{if } P_{\kappa(A)+i}(A) > 0, P_{i+k-\kappa(\neg A)}(\neg A) > 0 \\ P_{\kappa(A)+i}(\cdot|A) & \text{if } P_{\kappa(A)+i}(A) > 0, P_{i+k-\kappa(\neg A)}(\neg A) = 0 \\ P_{i+k-\kappa(\neg A)}(\cdot|\neg A) & \text{if } P_{\kappa(A)+i}(A) = 0, P_{i+k-\kappa(\neg A)}(\neg A) > 0 \end{cases}$$

¹⁴In particular, we only consider revision by consistent A (see also [22]). Gärdenfors permits revision by some inconsistent A to result in the inconsistent p-function P^\perp .

Theorem 4.4 *Let κ^P be any POCF that induces the ordered family $\{P_i\}$. Then the regular A, k -conditionalization of κ^P using mixture factors $\delta_0, \dots, \delta_m$ induces the A, k -conditionalized family of p -functions $\{P_i^{A,k}\}$ using the same mixture factors.*

Proof Recall that $P_i^{A,k}$ is defined (w.r.t. $\kappa_{A,k}^P$) using the weighting function restricted to $\kappa_{A,k}^{-1}(i)$. Assume $\kappa_{A,k}^{-1}(i)$ is an A -cluster. Then $\kappa^{-1}(i+k-\kappa(\neg A)) \cap \|\neg A\| = \emptyset$ and

$$P_i^{A,k}(B) = \frac{\sum\{P(w) : \kappa(w) = i + \kappa(A) \text{ and } w \models A \wedge B\}}{\sum\{P(w) : \kappa(w) = i + \kappa(A) \text{ and } w \models A\}}$$

Thus, $P_i^{A,k}(B) = P_{i+\kappa(A)}(B|A)$. However, since $\kappa_{A,k}^{-1}(i)$ is an A -cluster, $P_{i+k-\kappa(\neg A)}(\neg A) = 0$ and the direct A, k -conditionalization of $\{P_i\}$ also defines $P_i^{A,k}(B) = P_{i+\kappa(A)}(B|A)$.

Similar arguments can be made for resulting $\neg A$ -clusters and mixed clusters. ■

Regular A, k -conditionalization of a POCF or its ordered family has a number of attractive properties. However, in order to use this procedure for revision of a POCF, or its associated Popper and p -functions, a fair bit of knowledge has to be provided. Because POCFs generalize OCFs, the evidence A has to be specified with a degree of strength (or entrenchment) k . In addition, a sequence of mixture factors δ_i must also be specified. Unfortunately, these mixture factors do not seem to have a natural interpretation in general.¹⁵ There are few hints in the formal structure of a POCF that might guide the selection of the appropriate δ_i , or even a single fixed δ for all mixtures. A lack of principles for the selection of mixture factors and intuitive interpretation of their meaning and function make regular A, k -conditionalization, and perhaps POCFs, less attractive as models of probabilistic revision.

One possible solution to this problem is to treat a POCF as a representation of some classical p -function \overline{P} , as suggested in Section 3.3, and treat update by A, k as Jeffrey conditioning (setting $\overline{P}(A) = p$ for some appropriate value p) on this p -function. One may be able to select an appropriate probability p based on the strength of evidence k ; suitable parameter values for A, k -conditionalization might then be suggested by the classical p -function $\overline{P}^{A,p}$ that results from classical Jeffrey conditioning on \overline{P} . Specifically, recall that if a POCF κ^P induces

¹⁵One exception pertains to $A, 0$ -conditionalization. Such a revision is best thought of as contraction as it ensures that both A and $\neg A$ are considered completely unsurprising. The revised POCF $\kappa_{A,0}^P$ induces a new absolute p -function $P^{A,0}(\cdot) = P^{A,0}(\cdot|\top)$ whose structure is influenced considerably by δ_0 . In particular, we have $P^{A,0}(A) = \delta_0$ and $P^{A,0}(\neg A) = (1 - \delta_0)$. Thus, regular $A, 0$ -conditionalization is a form of weighted contraction that might be viewed as some type Jeffrey conditionalization setting $P(A) = \delta_0$.

an ordered family $\{P_i : \kappa^{-1}(i) \neq \emptyset\}$, we require the underlying p-function \overline{P} to be some additive mixture of the P_i :

$$\overline{P}(A) = \sum \{a_i P_i(A) : \kappa^{-1}(i) \neq \emptyset\}$$

A, k -conditionalization determines a new POCF $\kappa_{A,k}^P$ which in turn determines a new ordered family $\{P_i^{A,k} : \kappa_{A,k}^{-1}(i) \neq \emptyset\}$. To interpret A, k -conditionalization as Jeffrey conditioning on \overline{P} , we require that some p-function $\overline{P}^{A,p}$ — determined by Jeffrey conditioning of \overline{P} — be equivalent to an additive mixture of this new ordered family:

$$\overline{P}^{A,p}(A) = \overline{P}^{A,k}(A) \equiv_{\text{df}} \sum \{b_i P_i^{A,k}(A) : \kappa_{A,k}^{-1}(i) \neq \emptyset\}$$

If such a value p can be found (along with appropriate additive weights b_i), it may be the case that appropriate mixture factors δ_i are fixed by this relationship.

It turns out that *if* any A, k -conditionalization can be interpreted as Jeffrey conditioning on an underlying classical p-function in the manner suggested, it must be (equivalent to) some *regular* A, k -conditionalization, as indicated by the following lemma. To prove this, we exploit the fact that the Jeffrey conditioning of P by $P(A) = p$ leaves the conditional p-functions $P(\cdot|A)$ and $P(\cdot|\neg A)$ unchanged. Indeed, to guarantee that P' can be formed using Jeffrey conditioning on A it is sufficient to show that $P'(C|A) = P(C|A)$ and $P'(C|\neg A) = P(C|\neg A)$ for all C .

Lemma 4.5 *Let κ^P be a POCF with ordered family $\{P_i\}$, and let \overline{P} be an additive mixture of $\{P_i\}$. Let $\kappa_{A,k}^P$ be some arbitrary A, k -conditionalization of κ^P with induced ordered family $\{P_i^{A,k}\}$, and let $\overline{P}^{A,k}$ be defined as some additive mixture of $\{P_i^{A,k}\}$. If $\overline{P}^{A,k} = \overline{P}^{A,p}$ (i.e., the Jeffrey conditioning of \overline{P} setting $P(A)=p$) for some p , then $\{P_i^{A,k}\}$ is a regular A, k -conditionalization of the family $\{P_i\}$.*

Proof We prove that the resulting mixed clusters of the A, k -conditionalization must conform to the constraints of regular A, k -conditionalization. The proof for A and $\neg A$ -clusters is similar.

Let $\kappa_{A,k}^{-1}(i)$ be some mixed cluster in $\kappa_{A,k}^P$ such that

$$P_i^{A,k} \neq \delta_i P_{\kappa(A)+i}(\cdot|A) + (1 - \delta_i) P_{i+k-\kappa(\neg A)}(\cdot|\neg A)$$

for any $0 < \delta_i < 1$. This implies either that a) for some pair of worlds $w, v \in \|A\|$ such that $\kappa_{A,k}(w) = \kappa_{A,k}(v) = i$, we have

$$\frac{P_i^{A,k}(w)}{P_i^{A,k}(v)} \neq \frac{P_{\kappa(A)+i}(w)}{P_{\kappa(A)+i}(v)}$$

or b) a similar relation for the $\neg A$ -part of $\kappa_{A,k}^{-1}(i)$. (We assume the former; the proof of the latter case is identical.)

Since \overline{P} is an additive mixture of $\{P_i\}$ and $\overline{P}^{A,k}$ is an additive mixture of $\{P_i^{A,k}\}$, we must have

$$\frac{\overline{P}^{A,k}(w)}{\overline{P}^{A,k}(v)} \neq \frac{\overline{P}(w)}{\overline{P}(v)}$$

Thus, $\overline{P}^{A,k}$ cannot be formed by Jeffrey conditioning on \overline{P} . ■

The Jeffrey conditioning of an underlying classical p-function justifies the use of regular A, k -conditionalization of κ^P ; but it still cannot aid in the selection of appropriate mixture factors δ_i . Again, assume that κ^P induces an ordered family $\{P_i\}$ and that $\overline{P} = \sum\{a_i P_i\}$ is an additive mixture of this family. If $\kappa_{A,k}^P$ induces the family $\{P_i^{A,k}\}$, we want $\overline{P}^{A,k} = \sum\{b_i P_i^{A,k}\}$ to be an additive mixture of this family. If $\overline{P}^{A,k}$ can also be produced by Jeffrey conditioning of \overline{P} , using Lemma 4.5 it is easy to verify the following fact.¹⁶

Lemma 4.6 $\overline{P}^{A,k}$ is produced by Jeffrey conditioning of \overline{P} by A iff 1) $\{P_i^{A,k}\}$ is the regular A, k -conditionalization of $\{P_i\}$ (using some factors δ_i); and 2) for any two clusters $\kappa_{A,k}^{-1}(i)$ and $\kappa_{A,k}^{-1}(j)$ of $\kappa_{A,k}^P$, the new additive weights b_i satisfy the following properties:

(a) If both are mixed clusters then

$$\frac{a_{\kappa(A)+i} \cdot P_{\kappa(A)+i}(A)}{a_{\kappa(A)+j} \cdot P_{\kappa(A)+j}(A)} = \frac{\delta_i b_i}{\delta_j b_j}$$

and

$$\frac{a_{i+k-\kappa(\neg A)} \cdot P_{i+k-\kappa(\neg A)}(\neg A)}{a_{j+k-\kappa(\neg A)} \cdot P_{j+k-\kappa(\neg A)}(\neg A)} = \frac{(1-\delta_i)b_i}{(1-\delta_j)b_j}$$

(b) If $\kappa_{A,k}^{-1}(i)$ is a mixed cluster and $\kappa_{A,k}^{-1}(j)$ is an A -cluster then

$$\frac{a_{\kappa(A)+i} \cdot P_{\kappa(A)+i}(A)}{a_{\kappa(A)+j} \cdot P_{\kappa(A)+j}(A)} = \frac{\delta_i b_i}{b_j}$$

(c) If $\kappa_{A,k}^{-1}(i)$ is a mixed cluster and $\kappa_{A,k}^{-1}(j)$ is a $\neg A$ -cluster then

$$\frac{a_{i+k-\kappa(\neg A)} \cdot P_{i+k-\kappa(\neg A)}(\neg A)}{a_{j+k-\kappa(\neg A)} \cdot P_{j+k-\kappa(\neg A)}(\neg A)} = \frac{(1-\delta_i)b_i}{b_j}$$

¹⁶A proof is based on the same considerations as those used in the proof of the lemma, namely, that the conditional p-functions $P(\cdot|A)$ and $P(\cdot|\neg A)$ must remain unchanged.

If the additive weights b_i are varied, the constraints on suitable mixture factors δ_i also change. Rather than restricting the choice of appropriate δ_i , we have introduced yet another parameter to be tuned in the revision of a p-function. This is not so problematic, for we may assume that the additive weights are fixed. After any revision of a POCF, the weight assigned to each p-function $P_i^{A,k}$ can be simply that assigned to its predecessor P_i ; that is, $b_i = a_i$ in the above scheme. Of course, this is too stringent, for the number of nontrivial clusters or p-functions may change in the move from κ^P to $\kappa_{A,k}^P$. More realistically, we might imagine that these weights be fixed, but normalized to discount those weights a_i that correspond to empty clusters. So suppose $\overline{P} = \sum \{a_i P_i\}$ and $\overline{P}^{A,k} = \sum \{b_i P_i^{A,k}\}$. We call the new p-function $\overline{P}^{A,k}$ a *proportional* revision of \overline{P} (w.r.t. to κ^P) just when $\frac{a_i}{a_j} = \frac{b_i}{b_j}$ whenever the clusters $\kappa^{-1}(i)$, $\kappa^{-1}(j)$, $\kappa_{A,k}^{-1}(i)$ and $\kappa_{A,k}^{-1}(j)$ are nonempty. Unfortunately, the requirement of proportionality conflicts with the interpretation of A, k -conditionalization as Jeffrey conditioning.

Theorem 4.7 *There exists a POCF κ_P such that for no additive mixture \overline{P} of its ordered family is there a proportional additive mixture $\overline{P}^{A,k}$ of its A, k -conditionalization that can be constructed from Jeffrey conditioning of \overline{P} .*

We present a simple counterexample to verify this fact. Let κ^P be a POCF with four clusters and associated ordered family of p-functions $\{P_0, P_1, P_2, P_3\}$. Assume $P_0(A) = P_1(A) = 1$ and $P_2(\neg A) = P_3(\neg A) = 1$ and that $\overline{P} = a_0 P_0 + a_1 P_1 + a_2 P_2 + a_3 P_3$. The $A, 1$ -conditionalization of κ^P causes the $\neg A$ -worlds to be shifted down one level; thus $\kappa_{A,1}^P$ consists of three clusters. By Lemma 4.5, if $\overline{P}^{A,1}$ is to be equivalent to the Jeffrey conditionalization of \overline{P} , then the ordered family induced by $\kappa_{A,1}^P$ must have the form, for some $0 < \delta_1 < 1$: $P_0^{A,1} = P_0$; $P_1^{A,1} = \delta_1 P_1 + (1 - \delta_1) P_2$; and $P_2^{A,1} = P_3$. Now, let $\overline{P}^{A,1} = b_0 P_0^{A,1} + b_1 P_1^{A,1} + b_2 P_2^{A,1}$. By the constraints described above, if $\overline{P}^{A,1}$ is to be equivalent to the Jeffrey conditionalization, we have

$$\frac{a_0}{a_1} = \frac{b_0}{\delta_1 b_1}$$

However, if $\overline{P}^{A,1}$ is a proportional revision of \overline{P} then by Lemma 4.6 we must also have

$$\frac{a_0}{a_1} = \frac{b_0}{b_1}$$

This contradicts the fact that $0 < \delta_1 < 1$.

We notice that this counterexample is a very typical form of probability revision, and most such “run of the mill” revisions will give rise to the same “impossibility” result. Thus, to treat POCFs as abstractions of absolute p-functions in order to determine mixture weights δ_i by appeal to Jeffrey conditioning, one

must propose criteria according to which new additive factors b_i should be selected. Thus, rather than restricting the choice of δ_i , one introduces yet another choice, another parameter that must be fixed in the updating of POCF. Revision of Popper functions using POCFs is somewhat unattractive because of the epistemological demands on the holder of an epistemic state and the provider of evidence by which the epistemic state is to be updated. For this reason, it is certainly worthwhile exploring simpler alternatives.

4.2 Probabilistic Minimal Conditional Revision

Spohn’s revision method allows for iteration, but cannot be applied to QRMs, for the relative plausibility of worlds is not sufficient to determine an updated ranking; the actual κ -ranking of a world’s plausibility is necessary. For this reason, the extension of Spohn’s method to POCFs cannot be applied to Popper functions directly — the ordered family representation of a Popper function merely determines the relative plausibility of worlds and p-functions, not the magnitude of plausibility. Methods of iterated revision that work directly with QRMs are therefore most directly applicable to the problem of Popper function revision.

Several such proposals have been put forth. Safe contraction [3], generalized epistemic entrenchment [25] and the probabilistically motivated system of Schlechta [26] each take a similar approach to the problem: each assumes the existence of a “global” ordering of entrenchment over all sentences in the language. For any belief set K the appropriate revision function is immediately available, and iteration of the process requires no additional apparatus. These models have the rather severe drawback that any objective belief set K is associated with a unique revision function. Furthermore, such an ordering determines globally preferred belief sets. Hansson [16] proposes that instead a revision method be associated with *belief bases* rather belief sets. Thus, the same belief set may be revised in different ways if it is generated by different bases in each instance. In our setting, the revision of a belief set need not be tied to its underlying belief base.¹⁷

We will examine in detail the probabilistic extension of the method of *minimal conditional (MC) revision* [7, 5, 4]. Given a QRM \leq suitable for some belief set K , a scheme for iterated revision must produce not only a revised belief set K_A^* , but also a new QRM \leq_A^* . This QRM must satisfy the Basic Requirement that $\min(\leq_A^*, \top) = \min(\leq, A)$. MC-revision is based on the intuition that the rest of the structure of \leq should be left intact to the greatest extent possible. The set of worlds $\min(\leq, A)$ becomes most plausible, while the relative plausibility of all other worlds remains unchanged.

Definition 4.6 Let \leq be a QRM. The *MC-revision operator* $*$ maps \leq into

¹⁷Of course, such ties can be added. Work on base contraction [15, 16, 11] can be viewed in this light.

\leq_A^* , for any $A \in \mathbf{L}$, where a) if $v \in \min(\leq, A)$ then $v \leq_A^* w$ for all $w \in W$ and $w \leq_A^* v$ iff $w \in \min(\leq, A)$; and b) if $v, w \notin \min(\leq, A)$ then $w \leq_A^* v$ iff $w \leq v$.

Given \leq_A^* , the revised belief set K_A^* is defined in the obvious way as those sentences true on the set $\min(\leq_A^*, \top)$, clearly satisfying the AGM postulates. Iterated revision of K_A^* proceeds using the new QRM \leq_A^* to guide the process.

We do not elaborate on the properties of MC-revision here. However, we note that this method produces a new QRM that preserves as much of the original entrenchment relation as is consistent with the AGM postulates.¹⁸ Also of interest is the fact that any sequence of revisions A_1, \dots, A_n can be reduced to a single uniterated revision A ; that is, there is a characteristic sentence $A \in \mathbf{L}$ for the sequence such that $((K_{A_1}^*)_{A_2}^* \dots)_{A_n}^* = K_A^*$. Furthermore, this A can be determined using the entrenchment information captured by the original QRM \leq . However, the ordering $((\leq_{A_1}^*)_{A_2}^* \dots)_{A_n}^*$ is generally not equivalent to \leq_A^* (nor generally does there exist a single A that has the same effect on the ordering).

The method of MC-revision can be extended to CPMs, the probabilistic counterpart of QRMs, in a rather straightforward way, just as A, k -conditionalization was extended to POCFs.

Definition 4.7 Let $\leq^P = \langle \leq, P \rangle$ be a CPM. The MC-revision of \leq^P by (consistent) $A \in \mathbf{L}$ is $\leq_A^P = \langle \leq_A^*, P \rangle$, where \leq_A^* is the MC-revision of \leq by A .

Note that the weighting function P remains unchanged in the move from \leq^P to \leq_A^P . This is feasible because of the structure of MC-revision: the cluster $\min(\leq, A)$ is “split” and its A -part becomes most plausible, while all other clusters remain unchanged. Thus, no clusters are combined and the relative weights need not be altered to preserve the appropriate conditional probabilities (given A and $\neg A$). This stands in stark contrast with the potentially drastic changes to weights required when revising POCFs. In particular, we have the following.

Theorem 4.8 Let \leq^P be a CPM and $P(\cdot \uparrow \cdot)$ its corresponding Popper function. Let \leq_A^P be the MC-revision of \leq^P and $P^A(\cdot \uparrow \cdot)$ its corresponding Popper function. Then $P(B \uparrow A \wedge C) = P^A(B \uparrow A \wedge C)$ and $P(B \uparrow \neg A \wedge C) = P^A(B \uparrow \neg A \wedge C)$ for all B, C .

Proof Assume $A \wedge C$ is satisfiable (otherwise $P(B \uparrow A \wedge C) = P^{A,k}(B \uparrow A \wedge C) = 1$ for all B). We have

$$P(B \uparrow A \wedge C) = \frac{\sum \{P(w) : w \in \min(\leq, A \wedge C) \text{ and } w \models B\}}{\sum \{P(w) : w \in \min(\leq, A \wedge C)\}}$$

¹⁸We refer to [7, 4] for details; equivalently, MC-revision produces the minimal possible change in an agent’s set of *conditional beliefs*.

We also have $\min(\leq_A^P, A \wedge C) = \min(\leq, A \wedge C)$. Since the weights of the worlds in $\min(\leq, A \wedge C)$ is unchanged, $P(B \uparrow A \wedge C) = P^A(B \uparrow A \wedge C)$. An analogous argument can be made for $P^A(B \uparrow \neg A \wedge C)$. ■

Corollary 4.9 $P(B \uparrow A) = P^A(B \uparrow A)$ and $P(B \uparrow \neg A) = P^A(B \uparrow \neg A)$.

Corollary 4.10 $P^A(B \uparrow \top) = P(B \uparrow A)$.

It is easy to see, as a result, that MC-revision of \leq^P determines a probability revision operation satisfying postulates (P1)–(P5) if we take P to be the absolute p-function induced by \leq^P and P_A^* to be the p-function induced by \leq_A^P .

We can apply the operation of MC-revision directly to the minimal, orthogonal ordered family of p-functions induced by a CPM. Intuitively, a family P_0, \dots, P_n mutates into the sequence $P_k(\cdot|A), P_0, \dots, P_{k-1}, P_k(\cdot|\neg A), P_{k+1}, \dots, P_n$, where P_k is the first p-function in the original sequence that gives A positive probability. Of course, if $P_k(\neg A) = 0$, the term $P_k(\cdot|\neg A)$ is deleted from the sequence.

Definition 4.8 Let $\{P_i : 0 \leq i \leq n\}$ be the ordered family of p-functions induced by some CPM \leq^P . Let P_k be the minimal A -permitting p-function in this sequence for some consistent $A \in \mathbf{L}$. The MC-revised family of p-functions $\{P_i^A\}$ is defined as follows:

$$\begin{aligned} P_0^A &= P_k(\cdot|A) \\ P_{i+1}^A &= P_i && \text{for } 0 \leq i < k \\ P_{k+1}^A &= P_k(\cdot|\neg A) && \text{if } P_k(\neg A) > 0 \\ &= P_{k+1} && \text{if } P_k(\neg A) = 0 \\ P_{i+1}^A &= P_i && \text{for } k < i \leq n, \text{ if } P_k(\neg A) > 0 \\ &= P_{i+1} && \text{for } k < i \leq n, \text{ if } P_k(\neg A) = 0 \end{aligned}$$

The following theorem is immediate.

Theorem 4.11 Let \leq^P be any CPM that induces the ordered family $\{P_i\}$. Then the MC-revision of \leq^P by A induces the MC-revised family of p-functions $\{P_i^A\}$.

There are some crucial differences between MC-revision and regular A, k -conditionalization. First, since clusters in a CPM can only be split by MC-revision, the need for mixture factors (the δ_i used above) is obviated. Furthermore, revision by A need not be accompanied by a degree of entrenchment or “weight of evidence” parameter k as is the case for A, k -conditionalization. Finally, since the result of MC-revision is determined solely by the structure of a CPM or its ordered family representation (and not by the magnitudes of plausibility measures), it is well-defined for any Popper function and uniquely determines a revised Popper function. In general, we will take the MC-revision

of a Popper function to be the Popper function corresponding to the MC-revision of its minimal, orthogonal ordered family representation.

One drawback of MC-revision is that new beliefs are accepted with what might be termed a “minimal” degree of entrenchment. Only the most plausible A -worlds are shifted in relative plausibility; if some new fact B is learned subsequently and $\neg B \in K_A^*$ then A is at great risk of being retracted.¹⁹ However, without degrees of entrenchment whose magnitudes can be compared, this might be the best we can hope for. Furthermore, the use of MC-revision to revise Popper functions has the following appealing property: subject to the constraints of (P1)–(P5), the MC-revision of a Popper function changes as few conditional probabilities as possible. In other words, MC-revision represents the minimal possible change of a Popper function required to capture revision by A .²⁰

Lemma 4.12 *Let $P(\cdot \uparrow \cdot)$ be a Popper function and let $P^A(\cdot \uparrow \cdot)$ be the Popper function determined by the MC-revision of P by any consistent $A \in \mathbf{L}$. If $P(B \uparrow A) = 0$ then $P^A(C \uparrow B) = P(C \uparrow B)$ for all $B, C \in \mathbf{L}$.*

Proof We assume B, C are consistent, for the lemma holds trivially otherwise.

Let $\{P_i\}$ be the ordered family of p-functions induced by P and let $\{P_i^A\}$ be the revised ordered family determining P^A . Denote by P_C (resp. P_C^A) the minimal C -permitting p-function in $\{P_i\}$ (resp. $\{P_i^A\}$). Since $P(B \uparrow A) = 0$, we have that P_B and P_A are distinct. By definition of $\{P_i^A\}$, we are guaranteed that $P_B = P_B^A$; thus, $P^A(C \uparrow B) = P(C \uparrow B)$ for all B, C . ■

Lemma 4.13 *Let $P(\cdot \uparrow \cdot)$ be a Popper function and let $P^A(\cdot \uparrow \cdot)$ be the Popper function determined by the MC-revision of P by any consistent $A \in \mathbf{L}$. If $P(B \uparrow A) > 0$ then $P^A(C \uparrow B) = P(C \uparrow A \wedge B)$ for all $B, C \in \mathbf{L}$.*

Proof The proof proceeds as that of the previous lemma. By definition of $\{P_i^A\}$, we have that $P_0^A = P_A(\cdot | A)$. Since $P(B \uparrow A) > 0$, we have $P_A(B | A) = P_0^A(B) > 0$. Thus $P_B^A = P_0^A$ and $P^A(\cdot \uparrow B) = P_0^A(\cdot | B) = P_A(\cdot | A \wedge B)$. So, we have $P^A(C \uparrow B) = P(C \uparrow A \wedge B)$ for all B, C . ■

¹⁹A similar point is made by Spohn [28] who briefly describes and dismisses a proposal much like MC-revision for QRMs.

²⁰By “minimal change” we mean that $P^A(C \uparrow B) = P(C \uparrow B)$ for as many C, B as possible, where P^A is the Popper function induced by MC-revision of P . Other notions of minimal change, as applied to p-functions, include the cross-entropy measure of a distribution and its revised counterpart [32, 31]. Unfortunately, such a measure is not directly applicable to Popper functions — unless we give them a classical interpretation as in Section 3. In the case where such a measure is applicable — when $P(A \uparrow \top) > 0$ — both MC-revision and A, k -conditionalization hold up to the test with respect to the absolute p-function $P(\cdot \uparrow \top) > 0$, for both perform ordinary conditioning by A on this p-function, and therefore minimize cross-entropy.

Notice that the conditional probability $P^A(C \uparrow B)$ may be different from $P(C \uparrow B)$ after revision in the case where $P(B \uparrow A) > 0$, as indicated by Lemma 4.13. However, these changes are required if the revision function applied to the induced absolute function $P(\cdot|\top)$ is to satisfy postulate (P5). By Lemma 4.12 all other conditional probabilities are unchanged. Thus, we are guaranteed that MC-revision minimally changes the Popper function. Formally, we say that P' is more *similar* to P than P'' is just when

$$\{\langle A, B \rangle : P''(A \uparrow B) = P(A \uparrow B)\} \subset \{\langle A, B \rangle : P'(A \uparrow B) = P(A \uparrow B)\}$$

Theorem 4.14 *Let $*$ be a revision function satisfying (P1)–(P5). Let $P(\cdot \uparrow \cdot)$ be a Popper function with underlying p-function $P(\cdot|\top)$, and let $P'(\cdot \uparrow \cdot)$ be a Popper function suitable for P_A^* . Then P' is maximally similar to P iff P' is the MC-revision of P .*

Finally, we see that the information content of the sequence of underlying p-functions induced by a sequence of revisions is nondecreasing.

Definition 4.9 Let P, Q be p-functions. We say P is *less informative* than Q iff Q can be obtained from P by nontrivial conditionalization; that is, if $Q = P(\cdot|A)$ for some A such that $0 < P(A) < 1$.

Proposition 4.15 *Let $P(\cdot \uparrow \cdot)$ be a Popper function and $P = P(\cdot \uparrow \top)$ its underlying p-function. Let $P^{A_1}, (P^{A_1})^{A_2}, \dots, (((P^{A_1})^{A_2}) \dots)^{A_n}$ be the sequence of p-functions induced by the revision of P by A_1, \dots, A_n . Then $(((P^{A_1})^{A_2}) \dots)^{A_1}$ is not less informative than $(((P^{A_1})^{A_2}) \dots)^{A_j}$ if $i \geq j$.*

5 Concluding Remarks

We have presented some considerations on the iterated revision of probability functions. We have described two possible models of the process. The first is based on a probabilistic extension of Spohn's OCFs and updating mechanism. Difficulties arise due to the epistemological demands placed on the epistemic state of POCFs and on the provider of evidence. It remains to be seen if reasonable criteria can be proposed for the selection of mixture factors required to combine p-functions in the manner dictated by Spohn's proposal. A second model, MC-revision, is based more directly on the structural properties of Popper functions and allows for minimal changes in an agent's conditional probabilities. Unfortunately, this model does not allow degrees of entrenchment to be associated with evidence (nor could it deal with those if they were provided). As such, the minimal and *weakest* change to the Popper function is adopted. We conclude that the revision of probabilistic belief states is not as well-understood as we might have imagined, and that it is not as well-behaved as we might hope.

Acknowledgements

I would like to thank Wolfgang Spohn for valuable comments on an earlier draft of this paper, and Nir Friedman and Adam Grove for helpful pointers to relevant literature. This research was supported by NSERC Research Grant OGP0121843.

References

- [1] Ernest W. Adams. *The Logic of Conditionals*. D.Reidel, Dordrecht, 1975.
- [2] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [3] Carlos Alchourrón and David Makinson. On the logic of theory change: Safe contraction. *Studia Logica*, 44(4):405–422, 1985.
- [4] Craig Boutilier. Sequences of revisions: On the semantics of nested conditionals. Technical Report 92-24, University of British Columbia, Vancouver, 1992.
- [5] Craig Boutilier. Revision sequences and nested conditionals. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 519–525, Chambery, FR, 1993.
- [6] Craig Boutilier. Conditional logics of normality: A modal approach. *Artificial Intelligence*, 68:87–154, 1994.
- [7] Craig Boutilier. Iterated revision and minimal revision of conditional beliefs. *Journal of Philosophical Logic*, 1994. (in press).
- [8] Craig Boutilier. Toward a logic for qualitative decision theory. In *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning*, pages 75–86, Bonn, 1994.
- [9] Craig Boutilier. Unifying default reasoning and belief revision in a modal framework. *Artificial Intelligence*, 68:33–85, 1994.
- [10] Peter Cheeseman. In defense of probability. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 1002–1009, Los Angeles, 1985.
- [11] André Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20:175–203, 1991.
- [12] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, 1988.

- [13] Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the Third Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 83–95, Pacific Grove, 1988.
- [14] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
- [15] Sven Ove Hansson. In defense of base contraction. *Synthese*, 91:239–245, 1992.
- [16] Sven Ove Hansson. In defense of the Ramsey test. *Journal of Philosophy*, 89:522–540, 1992.
- [17] Gilbert Harman. *Change in View*. MIT Press, Cambridge, 1986.
- [18] William L. Harper. Rational belief change, Popper functions and counterfactuals. In W.L. Harper and C.A. Hooker, editors, *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, volume 1, pages 73–115. D. Reidel, Dordrecht, 1976.
- [19] Richard C. Jeffrey. *The Logic of Decision*. McGraw-Hill, New York, 1965.
- [20] Henry E. Kyburg, Jr. *Probability and the Logic of Rational Belief*. Wesleyan University Press, Middletown, 1961.
- [21] Isaac Levi. *The Enterprise of Knowledge*. MIT Press, Cambridge, 1980.
- [22] Sten Lindström and Włodzimierz Rabinowicz. On the probabilistic representation of non-probabilistic belief revision. *Journal of Philosophical Logic*, 18:69–101, 1989.
- [23] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, 1988.
- [24] David Poole. Decision-theoretic defaults. In *Proceedings of Canadian Society for Computational Studies of Intelligence Conference*, pages 190–197, Vancouver, 1992.
- [25] Hans Rott. Belief change using generalized epistemic entrenchment. *Journal of Logic, Language and Information*, 1(1):45–78, 1992.
- [26] Karl Schlechta. Theory revision and probability. *Notre Dame Journal of Formal Logic*, 307(2):307–318, 1991.
- [27] Wolfgang Spohn. The representation of popper measures. *Topoi*, 5:69–74, 1986.

- [28] Wolfgang 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*, volume 2, pages 105–134. D. Reidel, Dordrecht, 1987.
- [29] Robert C. Stalnaker. Probability and conditionals. In W.L. Harper, R. Stalnaker, and G. Pearce, editors, *Ifs*, pages 107–128. D. Reidel, Dordrecht, 1970. 1981.
- [30] Bas C. van Fraassen. Representation of conditional probabilities. *Journal of Philosophical Logic*, 5:417–430, 1976.
- [31] Bas C. van Fraassen. Rational belief and probability kinematics. *Philosophy of Science*, 47:165–187, 1980.
- [32] P. M. Williams. Bayesian conditionalization and the principle of minimum information. *British Journal for the Philosophy of Science*, 31:131–144, 1980.