Ranking Functions for Belief Change
A Uniform Approach to Belief Revision and Belief Progression

Aaron Hunter
British Columbia Institute of Technology, Burnaby, Canada

Keywords: Belief Revision, Reasoning about Action, Knowledge Representation.

Abstract: In this paper, we explore the use of ranking functions in reasoning about belief change. It is well-known that the semantics of belief revision can be defined either through total pre-orders or through ranking functions over states. While both approaches have similar expressive power with respect to single-shot belief revision, we argue that ranking functions provide distinct advantages at both the theoretical level and the practical level, particularly when actions are introduced. We demonstrate that belief revision induces a natural algebra over ranking functions, which treats belief states and observations in the same manner. When we introduce belief progression due to actions, we show that many natural domains can be easily represented with suitable ranking functions. Our formal framework uses ranking functions to represent belief revision and belief progression in a uniform manner; we demonstrate the power of our approach through formal results, as well as a series of natural problems in commonsense reasoning.

1 INTRODUCTION

The study of *belief revision* is concerned with the manner in which an agent's beliefs change in response to new information. Following the highly influential AGM model (Alchourrón et al., 1985), many approaches to belief revision rely on some form of *entrenchment ordering*. The idea is simple: an a priori ordering over states is used to guarantee that an agent always believes the formulas that are true in the "most entrenched" states consistent with an observation. The literature is not always clear on exactly what the ordering represents: in some cases, it may represent the likelihood of each state, whereas in other instances it may represent an agent's strength of belief. It is well known that AGM revision can also be framed in terms of ranking functions (Spohn, 1988). In this paper, we illustrate that ranking functions have significant advantages in modelling entrenchment, particularly when agents are able to execute state-changing actions. We present a uniform approach to modeling belief revision as well as *belief progression*, which is the change in belief that occurs when an action is executed. We illustrate through formal results and practical examples that there are many situations where the choice between ranking functions and entrenchment orderings is significant.

1.1 Motivation

Belief revision is often described as the belief change that occurs when an agent receives new information about a static world. For example, an agent might believe that the lamp is off in a certain room behind a closed door. If the door is opened to reveal the lamp is on, then the agent must modify their beliefs to incorporate this fact. One way to model this form of belief change is by assuming an underlying ordering $\prec$ over all possible states, where precedence is understood to represent plausibility. The $\prec$-minimal states would initially be ones in which the lamp was off. After observing the light is on, the agent will believe the actual state is among the least $\prec$-states that are consistent with this observation.

We are interested in domains where plausibility can not easily be captured by an ordering. For example, there are cases where evidence is *additive*; the agent might require two reports that the light is on before changing beliefs. In this case, an observation of light might make certain states more plausible, without actually changing the relative order of possible states. Similiarly, there are cases were observations are *graded*; light under the door could indicate the lamp is on, or perhaps that the window is open. Finally there are cases where actions may have *unlikely effects*; opening the door might accidentally turn the

Hunter A..
Ranking Functions for Belief Change - A Uniform Approach to Belief Revision and Belief Progression.
DOI: 10.5220/0004812704120419
In *Proceedings of the 6th International Conference on Agents and Artificial Intelligence* (ICAART-2014), pages 412-419
ISBN: 978-989-758-015-4

light on. Each of these is difficult to capture in a situation where uncertainty is represented by an ordering, without any mechanism for comparing magnitudes of unlikely effects. Our aim in this paper is to provide a flexible approach where uncertainty over states, actions, and observations is modeled by ranking functions that can be compared and combined with basic arithmetic.

1.2 Contributions

This paper makes several contributions to existing work on belief change. First, we introduce a natural algebra of belief revision that simplifies the semantics for the general case, while simultaneously subsuming AGM revision. Second, we introduce a formal model of belief progression due to actions that is Markovian, but still allows belief revision to respect the action history by excluding certain states after actions are executed. Finally, we demonstrate through a series of commonsense reasoning examples that a great deal of practical expressive power can be gained by allowing plausibility functions to range not only over states, but also over possible actions.

2 BACKGROUND

2.1 AGM Belief Revision

We focus on propositional belief revision, so we assume an underlying propositional signature. One of the most influential approaches to belief revision is the AGM approach (Alchourrón et al., 1985). In the AGM approach, a *belief set* is a deductively closed set of formulas. In this paper, we restrict attention to finite propositional signatures, so a belief set can be represented by a propositional formula ϕ. The new information to be incorporated is also represented by a single formula, say γ. An *AGM revision operator* is a binary function $*$ that satisfies the *AGM postulates*. The following reformulation of the postulates is due to Katsuno and Mendelzon (Katsuno and Mendelzon, 1992).

[R1] $\phi * \gamma$ implies γ.
[R2] If $\phi \wedge \gamma$ is satisfiable, then $\phi * \gamma \equiv \phi \wedge \gamma$.
[R3] If γ is satisfiable, then $\phi * \gamma$ is satisfiable.
[R4] If $\phi_1 \equiv \phi_2$ and $\gamma_1 \equiv \gamma_2$, then $\phi_1 * \gamma_1 \equiv \phi_2 * \gamma_2$.
[R5] $(\phi * \gamma) \wedge \beta$ implies $\phi * (\gamma \wedge \beta)$.
[R6] If $(\phi * \gamma) \wedge \beta$ is satisfiable, then $\phi * (\gamma \wedge \beta)$ implies $(\phi * \gamma) \wedge \beta$.

The class of AGM revision operators can be characterized in terms of orderings on states, where a *state* is just an interpretation of the underlying propositional signature. Let f be a function that maps every propositional formula ϕ to a total pre-order $\preceq_\phi$ over states. We say that f is a *faithful assignment* if and only if

1. If $s_1, s_2 \models \phi$, then $s_1 =_\phi s_2$.
2. If $s_1 \models \phi$ and $s_2 \not\models \phi$, then $s_1 \prec_\phi s_2$,
3. If $\phi_1 \equiv \phi_2$, then $\preceq_{\phi_1} = \preceq_{\phi_2}$.

The following characterization result indicates that every AGM operator can be understood in terms of minimization over a faithful assignment.

Proposition 1. *(Katsuno and Mendelzon, 1992) A revision operator $*$ satisfies [R1]-[R6] just in case there is a faithful assignment that maps each ϕ to an ordering $\preceq_\phi$ such that*

$$s \models \phi * \gamma \iff s \text{ is a } \preceq_\phi -\text{minimal model of } \gamma.$$

A similar characterization can be given using Spohn's ordinal conditional functions (Spohn, 1988), which are functions mapping each state to an ordinal.

2.2 Belief Change Due to Actions

We assume set $\mathbf{A}$ of action symbols. The effects of actions are described by a *transition function* $f : S \times \mathbf{A} \rightarrow S$. Hence, a transition function takes a state and an action as arguments, then it returns a new state. Informally, the output is the state that results from executing the given action in the given state. We are primarily concerned with actions that have *deterministic* effects, though we also allow *non-deteministic* effects in some examples. Throughout this paper, we will let the lower case letter a (possibly with subscripts) range over actions. A propositional signature $\mathbf{F}$ together with a transition function f over S is called an *action signature*.

In the literature, *belief update* refers to the belief change that occurs when an agent receives information about a change in the state of the world (Katsuno and Mendelzon, 1991). The classic approach to belief update defines the operation on propositional formulas, just as belief revision is defined on propositional formulas. This can lead to ambiguity because it is not always clear when an agent should perform revision and when they shoud perform update (Lang, 2007).

We avoid this ambiguity by modeling belief change due to actions through the operation commonly called *belief progression*. Belief progression operators take a belief state and an *action* as input, and return a new belief state that is obtained by progressing each possible world in accordance with the effects of the action. We remark that belief change in our formal approach is Markovian in that the new beliefs are completely determined by the orginal beliefs as well as the event (action or observation) that has occured. The basic model is not able to represent a class Markov Decision Process because we do not incorporate any notion of likelihood on action effects; we address this by adding such a measure in one of the examples in §4.

3 PLAUSIBILITY FUNCTIONS

3.1 An Algebra for Belief Revision

We formally define plausibility functions as follows.

Definition 1. *Let X be a non-empty set. A plausibility function over X is a function $r : X \to \mathbf{N}$ such that $r(x) = 0$ for at least one $x \in X$.*

If r is a plausibility function and $r(x) \leq r(y)$, then we say that x is at least as plausible as y. Plausibility functions are similar to ordinal conditional functions (Spohn, 1988), except that the domain can be any arbitrary set and we restrict the range to the natural numbers. This definition is based on a similar concept introduced in (Hunter and Delgrande, 2006).

When r is a plausibility function over states, we can identify the minimal elements of r with the states currently believed possible. Let

$$Bel(r) = \{x \mid r(x) = 0\}.$$

The *degree of strength* of a plausibility function r is the least n such that $n = r(v)$ for some $v \notin Bel(r)$. Hence, the degree of strength is a measure of how difficult it would be for an agent to abandon the currently believed set of states.

We use plausibility functions to represent initial belief states, and also to represent new information for revision. Hence, revision in this context is just a binary operator on plausibility functions. Given any pair of plausibility functions r_1 and r_2, we can define a new function $r_1 + r_2$ such that $(r_1 + r_2)(x) = r_1(x) + r_2(x)$. Of course, the sum of two plausibility functions need not be a plausbility function; but we can obtain an equivalent plausibility function by normalizing.

Definition 2. *Let r_1 and r_2 be plausibility functions over X, and let m be the minimum value of $r_1 + r_2$. Then $r_1 * r_2$ is the function on X defined as follows:*

$$r_1 * r_2(x) = r_1(x) + r_2(x) - m.$$

It should be clear that $r_1 * r_2$ is a plausibility function, because it attains a minimum value of 0. We use the symbol $*$ for this operation, because it can be seen as a generalization of AGM belief revision. To make this explicit, we introduce a basic definition.

Definition 3. *A plausibility function r is* two-valued *iff the range of r is a set of size 2. If r is two-valued, we write $|r| = \{s \mid r(s) = 0\}$.*

A formula can be represented by a two-valued plausibility function. We remark also that every plausibility function defines a total pre-order. Hence, AGM belief revision can be captured by taking a plausibility function over states (the initial beliefs) and adding a two-valued plausibility function (the formula for revision).

The class of plausibility functions is clearly closed under $*$. We state some other basic properties.

Proposition 2. *The operator $*$ is associative. i.e. $(r_1 * r_2) * r_3 = r_1 * (r_2 * r_3)$.*

We remark that many approaches to iterated revision are not associative, so this result suggests that our model of revision does not align directly with work in this area. We accept this difference, as it has been argued that none of the existing appoaches to iterated revision are completely satisfactory(Stalnaker, 2009).

In the following propositions, let r_I be the plausibility function such that $r_I(x) = 0$ for all x. We refer to r_I as the *identity* function. As an initial belief state, the identity function represents ignorance. As information for revision, it represents a null observation.

Proposition 3. *$r_I * r = r * r_I = r$ for any plausibility function r.*

Proposition 4. *For any plausibility function r, there is a plausibility function r^{-1} such that $r * r^{-1} = r^{-1} * r = r_I$.*

In abstract algebra, any closed system with an operator satisfying Propositions 2, 3 and 4 is called a *group*. If the operator is also commutative, the system is called an *abelian group*.

Proposition 5. *The class of plausibility functions is an abelian group under $*$.*

The fact that $*$ defines an abelian group means that we can exploit all of the known results about groups to analyze the symmetries and structure of revision under this definition.

3.2 Adding Actions

In this section, we assume an underlying set of action symbols $\mathbf{A}$ as well as a transition function f that describes the effects of the actions in $\mathbf{A}$. A *belief progression* operator is a function that maps an initial belief state to a new belief state, given that some action has been executed. When actions are introduced, we need to account for the fact that certain states may not be possible following the execution of a particular action.

To address this issue, we introduce the notion of an *extended plausibility function*.

Definition 4. *An extended plausibility function over X is a function $r : X \to \mathbf{N} \cup \{\infty\}$ such that $r(x) = 0$ for at least one $x \in X$.*

We define ∞ to be larger than every number in $\mathbf{N}$. Morever, we define addition as follows

$$p + \infty = \infty + p = \infty \text{ for any } p \in \mathbf{N} \cup \{\infty\}.$$

For any plausibility function r, let

$$imp(r) = \{s \mid r(s) = \infty\}.$$

Hence, $imp(r)$ is the set of states that are "impossible" according to the function r. On the other hand, we will refer to the complement of $imp(r)$ as the set of "possible" states.

Proposition 6. *If r_1, r_2 are extended plausibility functions, then $imp(r_1 * r_2) = imp(r_1) \cup imp(r_2)$.*

It follows that the set of possible states does not change when we revise by a plausibility function r with $imp(r) = \emptyset$. So we can think of an extended plausibility function as a plausibility function together with a set I of *impossible* states that are assigned the plausibility ∞.

We are now able to define belief progression with respect to extended plausibility functions.

Definition 5. *Let $a \in \mathbf{A}$ with deterministic transition function f, and let r be an extended plausibility function over $\mathbf{A}$. Then $r \cdot a$ is the extended plausibility function:*

$$r \cdot a(s) = \begin{cases} \min(\{r(s') \mid f(a, s') = s\}) & \\ \infty & \textit{otherwise} \end{cases}$$

This definition says that the plausibility of each state following an action is obtained by progressing forward the effects of actions. Hence, for each state s with plausibility k, we say that $f(a, s) = k$. This just means that the plausibility of the state remains the same before and after the execution of the action a. However, since there may be more than one initial state s with the same outcome, we take the minimum possible value. We need to assign ∞ to some states is because some states are not possible after executing a. For example, in a deterministic world, all states where the door is open will be impossible after an agent performs a *closedoor* action. Therefore, $r \cdot a$ is the natural shifting of r by the effects of s for states in the range of a; for states that are not possible following a, the plausibility is defined to be ∞.

In many applications, it is important to know if an extended plausibility function is consistent with a sequence of actions. We let $\overline{a}$ denote a sequence of actions of indeterminate length, and we let $r \cdot \overline{a}$ denote the sequential progression of r by each element of $\overline{a}$.

Definition 6. *An extended plausibility function r is consistent with the action sequence $\overline{a}$ just in case $imp(r) = imp(r_I \cdot \overline{a})$.*

Hence, r is consistent with $\overline{a}$ if the execution of $\overline{a}$ from a state of total ignorance leads to the same collection of impossible states. Since we are primarily interested in extended plausibility functions that result from actions, we simply say that an extended plausibility function is *consistent* if it is consistent with some action sequence.

Proposition 7. *If the extended plausibility function r is consistent, then $r \cdot a$ and $r * r'$ are also consistent, for any action a any plausibility function r'.*

Hence, if every state is initially possible, then we need only be concerned with consistent functions. The following result says that the outcome of a sequence of actions is consistent provided all states were initially possible.

Proposition 8. *Let r_1, r_2 be plausibility functions and let $\overline{a}$ be a sequence of actions. The extended plausibility function $r_1 \cdot a * r_2$ is consistent with $\overline{a}$.*

A set of so-called "interaction properties" have been proposed to ensure that observations following actions are incorporated in a sensible manner (Hunter and Delgrande, 2011). The main postulate can be reformulated in our notation as the following:

$$\text{P5. } Bel(r \cdot a * r') \subseteq Bel(r_I \cdot a)$$

The postulate $P5$ asserts that, regardless of the initial belief state and the observation, the final belief state must be a possible outcome of the action a. It is straightforward to show that Proposition 8 entails than interaction property $P5$.

3.3 Plausibility Functions Over Actions

A plausibility function over action symbols can be used to represent an agent's uncertainty about the actions that have been executed. We let A range over plausibility functions over actions; so we think of A as a partially observed action. We extend the use of $\cdot$ in the following definition.

Definition 7. *Let r be an extended plausibility function and let A be a plausibility function over* $\mathbf{A}$. *Then $r \cdot A$ is the extended plausibility function such that:*

$$r \cdot A(s) = \min(\{r(s') + A(a) \mid f(a, s') = s\}).$$

Informally, the definition just says that the most likely final states are the states that result starting from the most likely initial states and carrying out the most likely actions.

Proposition 9. *Let A be an extended plausibility function over actions, and suppose that $Bel(A) = \{a\}$ and $A(a') = \infty$ for $a' \neq a$. Then for any extended plausibility function r, it follows that $r \cdot A = r \cdot a$.*

This proposition deals with the case where A essentially picks out a single action a. In this case, we get the same result we would get if we simply used progression by the action a.

The following proposition gives some indication of the role played by ∞ when we have uncertainty over actions.

Proposition 10. *If r is a plausibility function (i.e. with domain $\mathbf{N}$), then $r \cdot A(s) = \infty$ just in case one of the following holds:*

1. *A is inconsistent, or*
2. *If $f(a,s') = s$, then $A(a) = \infty$.*

Hence, if A is consistent, then the only states excluded by $r \cdot A$ are those that are not possible outcomes of any action that is possible according to A.

Note that the model proposed here is only appropriate for action domains where the effects of an action can not fail. We are using plausibility functions to rank the likelihood that an action has occurred; if an agent believes an action has occurred, then the effects of that action must hold. We will see in the examples, however, that it is possible to include nondeterministic and failed actions by adding some additional ranking functions for effects.

4 REPRESENTING NATURAL ACTION DOMAINS

In this section, we demonstrate how sequences of plausibility functions can be used to represent natural action domains. In terms of notation, we use $INIT$ to represent the initial plausibility function. We use A and O (possibly with subscripts) to represent plausibility functions over actions and states, respectively. We refer to O as an observation, as it is a ranking function on states that provides new information. To be clear, athough $INIT$, A and O are all plausibility functions, it may be the case that $INIT \cdot A * O$ is an extended plausibility function.

We now introduce a sequence of examples. In each case, we assume that the actions and observations are (simple) plausiblity functions, and the final belief state is an extended plausbility function where some states are excluded. The final plausiblity values can be obtained by minimizing the sum of all plausibility values over actions and states, restricting attention to sequences of actions that are actually possible in the underlying transition system. The agent therefore maintains a consistent representation of the plausibility of a world together with the effects of actions.

Example (Additive Evidence)**.** Bob believes that he turned the lamp off in his office, but he is not completely certain. As he is leaving the building, he talks first to Alice and then to Eve. If only Alice tells him his lamp is still on, then he will believe that she is mistaken. Similarly, if only Eve tells him his lamp is still on, then he will believe that she is mistaken. However, if both Alice and Eve tell Bob that his lamp is still on, then he will believe that it is in fact still on.

The action signature contains, among others, a propositional variable *LampOn* and an action symbol $TurnLampOff$. The underlying transition system defines the effects of turning the lamp off in the obvious manner. Let *ON* denote the set of states in which *LampOn* is true. The following plausibility functions describe this action domain.

1. $INIT(s) = 0$ if $s \in ON$, $INIT = 10$ otherwise
2. $A_1(a) = 0$ if $a = TurnLampOff$, $A_1(a) = 3$ otherwise
3. $O_1(s) = 0$ if $s \in ON$, $O_1(s) = 2$ otherwise
4. $A_2(a) = 0$ if $a = null$, $A_2(a) = 10$ otherwise
5. $O_2(s) = 0$ if $s \in ON$, $O_2(s) = 2$ otherwise

It is easy to verify that, under this representation, two observations of *ON* are required to make Bob believe that he did not turn the lamp off.

Example (Graded Evidence)**.** Bob receives a gift that he estimates to be worth \$7. He is curious about the price, so he tries to glance quickly at the receipt without anyone noticing. He believes that the receipt says the price is \$3. This is far too low, so Bob concludes that he must have mis-read the receipt. Since a “3” looks very similar to an “8”, he concludes that the price on the receipt must have been \$8.

To represent this example, it is useful to assume that the set of actions includes a distinguished action symbol *null* that is just the identity function over the set S of states. Define the plausibility function A_1 such that $A_1(null) = 0$ and $A_1(a) = 10$ for every non-null action a, because Bob believes that no actions have occurred. We assume that there are propositional variables $Cost1, Cost2, \ldots, Cost9$ interpreted to represent the cost of the gift. We define a plausibility function

$INIT$ representing Bob's initial beliefs.

$$INIT(w) = \begin{cases} 0 \text{ if } w = \{Cost7\} \\ 1 \text{ if } w = \{Cost6\} \text{ or } w = \{Cost8\} \\ 3 \text{ otherwise} \end{cases}$$

Note that Bob initially believes that the cost is \$7, but it is comparatively plausible that this cost is one dollar more or less. Finally, we define a plausibility function O_1 representing the observation of the receipt.

$$O_1(w) = \begin{cases} 0 \text{ if } w = \{Cost3\} \\ 1 \text{ if } w = \{Cost8\} \\ 3 \text{ otherwise} \end{cases}$$

Bob believes that the observed digit was most likely a "3", with the most plausible alternative being the visually similar digit "8".

Given these plausibility functions, the most plausible conclusion is that the actual price is \$8; this is the result obtained through minimization. In order to draw this conclusion, Bob needs graded evidence about states of the world and he needs to be able to weight this information against his initial beliefs.

The preceding examples illustrate that there are commonsense reasoning problems in which an agent needs to consider aggregate plausibilities over a sequence of actions and observations. Plausibility functions are well-suited for reasoning about such problems. Total pre-orders over states, on the other hand, are not. In the case of graded evidence, the important point is that we need to be able to distinguish between initial plausibilities and some measure of similarity between observations. This is easy to represent using ranking functions; it is more difficult to represent using orderings, as we need to introduce some mechanism for combining the "levels" of an ordering.

4.1 Non-deterministic and Failed Actions

In this section, we consider actions with non-deterministic effects, including actions that may fail. Let f be a non-deterministic transition function, so $f(a,s)$ is a *set of states* that represents the possible outcomes when action a is executed in state s. Given such a transition function along with a plausibility function over actions, it is not possible to give a clear categorical procedure for choosing the effects of each action in the most plausible world histories. This problem can be solved by following (Boutilier, 1995), and attaching a plausibility value to the possible effects of each action.

Definition 8. *An effect ranking function is a function* δ *that maps every action-state pair* (a,s) *to a plausibility function over* $f(a,s)$.

Informally, an effect ranking function gives the likelihood of each possible effect for each action. A *non-deterministic plausibility function* is a pair $\langle r, \delta \rangle$ where r is a plausibility function over actions and δ is an effect ranking function.

Example (Unlikely Action Effects)**.** Consider an action domain involving a single propositional variable *LampOn* indicating whether or not a certain lamp is turned on. There are two action symbols *Press* and *ThrowPaper* respectively representing the acts of pressing on the light switch, or throwing a ball of paper at the light switch. Informally, throwing a ball of paper at the light switch is not likely to turn on the lamp. But suppose that an agent has reason to believe that a piece of paper was thrown at the lamp and, moreover, the lamp has been turned on. Define A_1 so that *ThrowPaper* is the most likely action at time 1.

	λ	*Press*	*ThrowPaper*
A_1	10	1	0

Next define $INIT$ and O_1 so that initially the light is off, and then the light is on.

	$\emptyset$	*LightOn*
$INIT$	10	0
O_1	0	10

Finally, we define an effect ranking function δ capturing the fact that pressing is more likely to turn the light on. This ranking function says nothing about which action has actually occurred.

	$\emptyset$	$\{LightOn\}$
$\delta(Press, LightOn)$	0	10
$\delta(Press, \emptyset)$	1	2
$\delta(ThrowPaper, LightOn)$	0	10
$\delta(ThrowPaper, \emptyset)$	2	1

Introducing effect ranking functions makes the distinction between action occurrences and action effects explicit, which in turn gives a straightforward treatment of failed actions.

5 DISCUSSION

5.1 Prioritizing Plausibility Functions

In formalizing commonsense examples, we had a sequence of plausibility functions over states and actions. We executed a sequence of operations iteratively by minimization over plausibility values at each step. It is not clear if this is always appropriate, particularly if we know in advance that there will be several plausibility functions to combine.

This issue has been discussed for belief revision in (Delgrande et al., 2006), where it is suggested that sequences of observations should should first be combined through some form of prioritization, and then the combination should be used as input for revision. In our context, this could be accomplished by using a large scalar multiple to prioritize certain ranking functions, followed by summation to merge all information together. It may also be interesting to consider non-summation based aggregates; the important point is that using a different aggregate does not introduce a fundamental change to our framework.

We remark, however, that there is a distinction that is lost here. In particular, there is a difference between an action that fails to occur and an action that occurs, but fails to produce an expected effect. Consider an agent that tries to drop a glass on the ground to break it. One possible outcome is that the agent executes the drop action but it fails to occur; perhaps the glass sticks to the agent's hand. Alternatively, the glass could be successfully dropped without breaking. In our framework, both of these events are represented by a dropping action with the null effect. In some cases, this might not be appropriate.

5.2 Action Formalisms

The issues addressed in this paper have been addressed in related action formalisms. We have already mentioned related work that has been done in the context of transition systems (Hunter and Delgrande, 2006). Similar work has also been done in the Situation Calculus (SitCalc). The SitCalc is an action formalism based on first-order logic, summarized in (Levesque et al., 1998). While the original formalism does not incorporate any epistemic notions, knowledge and belief have been added in extended versions. The most relevant work for comparison with our approach is the framework for iterated belief change in (Shapiro et al., 2011). In order to reason about belief change, a ranked set of possible initial situations is introduced and this set is refined over time as an agent performs sensing actions.

On the surface, our work is distinguished from the work in the SitCalc in that we use a less expressive representation of action effects. By using transition systems, we hope to focus entirely on the role played by the relevant ranking functions in belief change. The representation of belief change in the SitCalc does not use rankings to measure an agents perception of the action that has been executed, nor does it attempt to merge multiple forms of uncertainty due to graded evidence and epistemic entrenchment. Of course this is not a limitation of the approach: the SitCalc is a very expressive formalism that can be used to capture such notions. For instance, recent work has provided a treatment of beliefs about failed actions in the SitCalc (Delgrande and Levesque, 2012). Our view is that it can be simpler to iron out the main issues with respect to believe change in a simple AGM-like framework first, before migrating the solutions to a sophisticated formal framework such as the SitCalc.

5.3 Dynamic Epistemic Logic

Dynamic Epistemic Logic (DEL) is a broad term that generally refers to formal models of changing knowledge and belief following in the tradition of (Baltag et al., 1998). For a complete discussion of work in this area, we refer the reader to the extensive introduction in (van Ditmarsch et al., 2007). Broadly, work on DEL is distinguished from the work presented here in that DEL is based on the use of Kripke structures to model knowledge and belief.

Recent work in DEL has incorporated key notions of plausibility from the AGM tradition, as well as notions of graded belief. For example, (Lorini, 2011) provides an interesting example of graded belief in DEL. This work is actually quite similar in spirit to ours, and is fueled by the same kind of commonsense reasoning examples. By using simple ranking functinos over sets, our hope is to highlight the significant aspects of belief change that need to be modeled before committing to the representation of belief that is embodied by a Kripke structure.

5.4 Reasoning with Ordinals

The plausibility functions used in this paper are really a variation of Spohn's ordinal conditional functions(OCFs). In this paper, we have taken the position that a single infinite value denoted by ∞ can be useful for representing impossiblity. Adding this notion of impossibility makes our model largely equivalent to work in possibilistic logic, where quantitative measures of likelihood are combined with a "necessity measure" of 0 (Dubois and Prade, 2004).

While our focus in this paper has been on reasoning with a single infinite value, we propose that there are situations where we actually want greater expressive power for discussing impossible states. For example, we may want to reason about such states hypothetically; in these situations, it can actually be useful to allow plausibility values to range over all ordinals.

We suggest, for example, that it may be useful for an agent to reason counterfactually in worlds that have plausibility starting from the ordinal ω. Consider a statement of the form "The present king of France is bald." We can reasonably expect an agent to revise their beliefs about the present king of France, even if they do not believe in his existence. For example, one might be told "All french monarchs just had a hair transplant." In this case, it could be counterfactually concluded that the present king of France is no longer bald. This kind of reasoning can be modelled by identifying each limit ordinal (such as ω) with a hypothetical world configuration. In this manner, ordinals of the form $\omega + i$ can then be used as plausibility values to represent events of varying degrees of implausibility within these hypothetical worlds. The ordering on limit ordinals indicates which hypetheticals are the most outlandish, but we are able to perform revision across each in a uniform manner. We leave this application of plausibility functions for future work.

6 CONCLUSIONS

In this paper, we have discussed the use of plausiblity functions for reasoning about belief change, with a particular focus on action domains. We have demonstrated that Spohn-style ranking functions can be used to define an algebra of belief change, which can then be extended to reason about belief progression due to actions. We then used the same kind of ranking functions to represent uncertainty over the actions that have been executed. Through commonsense examples we demonstrated that ranking functions are a flexible tool that can capture many different kinds of uncertainty. At a formal level, it has long been known that AGM revision can be defined in terms of ranking functions or total pre-orders; but there is a sense in which pre-orders are simpler, and they have tended to be more popular in the literature. Our examples give many situations that are easier to represent with ranking functions, because the gap between different levels of plausibility is important. In many cases, our uniform treatment of belief states and observations simplifies the algebra of belief change.

REFERENCES

Alchourrón, C., Gärdenfors, P., and Makinson, D. (1985). On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50(2):510–530.

Baltag, A., Moss, L., and Solecki, S. (1998). The logic of public announcements, common knowledge and private suspicions. In *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98)*, pages 43–56.

Boutilier, C. (1995). Generalized update: Belief change in dynamic settings. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI 1995)*, pages 1550–1556.

Delgrande, J., Dubois, D., and Lang, J. (2006). Iterated revision as prioritized merging. In *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning (KR2006)*.

Delgrande, J. and Levesque, H. (2012). Belief revision with sensing and fallible actions. In *Proceedings of the International Conference on Knowledge Representation and Reasoning (KR2012)*.

Dubois, D. and Prade, H. (2004). Possibilistic logic: a retrospective and prospective view. *Fuzzy Sets and Systems*, 144(1):3–23.

Hunter, A. and Delgrande, J. (2006). Belief change in the context of fallible actions and observations. In *Proceedings of the National Conference on Artificial Intelligence(AAAI06)*.

Hunter, A. and Delgrande, J. (2011). Iterated belief change due to actions and observations. *Journal of Artificial Intelligence Research*, 40:269–304.

Katsuno, H. and Mendelzon, A. (1991). On the difference between updating a knowledge base and revising it. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR 1991)*, pages 387–394.

Katsuno, H. and Mendelzon, A. (1992). Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(2):263–294.

Lang, J. (2007). Belief update revisited. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI07)*, pages 2517–2522.

Levesque, H., Pirri, F., and Reiter, R. (1998). Foundations for the situation calculus. *Linköping Electronic Articles in Computer and Information Science*, 3(18):1–18.

Lorini, E. (2011). A dynamic logic of knowledge, graded beliefs and graded goals and its application to emotion modelling. In *LORI*, pages 165–178.

Shapiro, S., Pagnucco, M., Lesperance, Y., and Levesque, H. (2011). Iterated belief change in the situation calculus. *Artificial Intelligence*, 175(1):165–192.

Spohn, W. (1988). Ordinal conditional functions. A dynamic theory of epistemic states. In Harper, W. and Skyrms, B., editors, *Causation in Decision, Belief Change, and Statistics, vol. II*, pages 105–134. Kluwer Academic Publishers.

Stalnaker, R. (2009). Iterated belief revision. *Erkenntnis*, 70(1-2):189–209.

van Ditmarsch, H., van der Hoek, W., and Kooi, B. (2007). *Dynamic Epistemic Logic*. Springer.