

Belief Change with Uncertain Action Histories

Aaron Hunter

*British Columbia Institute of Technology
Burnaby, BC, Canada*

AARON_HUNTER@BCIT.CA

James P. Delgrande

*Simon Fraser University
Burnaby, BC, Canada*

JIM@CS.SFU.CA

Abstract

We consider the iterated belief change that occurs following an alternating sequence of actions and observations. At each instant, an agent has beliefs about the actions that have occurred as well as beliefs about the resulting state of the world. We represent such problems by a sequence of ranking functions, so an agent assigns a quantitative plausibility value to every action and every state at each point in time. The resulting formalism is able to represent fallible belief, erroneous perception, exogenous actions, and failed actions. We illustrate that our framework is a generalization of several existing approaches to belief change, and it appropriately captures the non-elementary interaction between belief update and belief revision.

1. Introduction

Many formal approaches have been introduced for reasoning about belief change in the context of actions and observations (Jin & Thielscher, 2004; Shapiro, Pagnucco, Lesperance, & Levesque, 2011; Delgrande & Levesque, 2012). In general, the underlying assumption is that agents perform belief update following actions and belief revision following observations. Existing formalisms, for the most part, have treated actions and observations independently, with little explicit discussion about the interaction between the two. In this paper, we consider the belief change that occurs due to an alternating sequence of actions and observations. We are interested in action domains where an agent may have erroneous beliefs, both about the state of the world as well as the action history.

Let K denote the beliefs of an agent, given by a set of possible worlds. For $1 \leq i \leq n$, let A_i denote an action and let O_i denote an observation. Informally, we are interested in sequences of the form

$$K \diamond A_1 * O_1 \diamond \cdots \diamond A_n * O_n \quad (1)$$

where \diamond is an update operator and $*$ is a revision operator. Our interpretation of this expression is flexible in that the actions may be understood to represent actions executed by a particular agent, or they may be exogenous. Note that such sequences may contain conflicting information. For example, the observation O_n may not be possible following the actions A_1, \dots, A_n . In this case, there are two options.

1. Reject O_n .
2. Accept O_n , and modify A_1, \dots, A_n .

In order to determine which option is preferable for a specific problem, an agent needs to be able to compare the plausibility of O_n with the plausibility of each A_i .

Expressions of the form (1) have previously been addressed, under the assumption that ontic action histories are infallible and recent observations take precedence over older observations (Hunter & Delgrande, 2011). Clearly, there are action domains in which these assumptions are not reasonable. In this paper, we propose a more flexible approach in which actions and observations are both represented by Spohn-style ranking functions. We take this approach because it gives a uniform treatment of the plausibility of beliefs, observations, and actions. When presented with conflicting information, an agent can simply compare the relative plausibility of each action and observation; a uniform representation of all events makes it straightforward to determine the most plausible sequence. Moreover, with quantitative plausibility values, we can encode a variety of distinct scenarios by manipulating the magnitudes of the plausibility values for alternative events.

This paper makes several contributions to existing work on epistemic action effects. The main contribution is a formal mechanism for reasoning about incorrect or weakly held beliefs related to action histories. Existing formalisms are generally unable to compare the plausibility of an action occurrence with the plausibility of a state of the world. This is a problem because, in practical reasoning problems, agents are often put in a position where they can either believe a certain fact holds or they can believe that some action has occurred. By using the same formal tool to represent beliefs about actions and states, we explicitly address the manner in which prior action occurrences are postulated or retracted in response to new observations.

A second contribution of this work is a flexible treatment of *weak* or *unreliable* observations. While some revision operators incorporate each new observation, it is obvious that this is not a desirable feature in many reasoning domains. In many cases, it is preferable to discard an unreliable observation if it conflicts with current, strongly held beliefs. However, our approach does not simply allow an unreliable observation to be discarded. Since we use ranking functions to represent observations, each observation also includes plausible alternatives. As such, each observation can actually provide evidence of several different states, and to differing degrees. We will see that this is useful in representing observations of states with similar appearance. Similarly, ranking functions can give a natural representation of additive evidence, in the form of observations that must occur several times before changing an agent's beliefs.

We formulate all of our results in a simple transition system framework that makes our treatment of action effects explicit and easy to compare with more elaborate action formalisms. However, our fundamental approach to dealing with uncertainty does not actually *require* transition systems to be used for action effects. We do not intend for the methodology described here to compete with alternative formalisms for representing epistemic action effects; rather, we provide a high-level approach to dealing with uncertain action histories that focuses specifically on the interaction between actions and observations. The significant feature of this work is that it demonstrates how a single representation of plausibility can be used to model iterated beliefs about actions and states simultaneously; the basic approach could be employed in other action formalisms. At a practical level, we demonstrate the utility of our work by giving a series of examples involving belief change in which the relative weight given to actions and observations varies. We suggest that our

work advances existing work on reasoning about epistemic action effects in that we provide a flexible, elaboration tolerant approach for capturing several different kinds of belief change that can occur when there is uncertainty about action occurrences. Note that this paper is an extended version of the work presented by Hunter and Delgrande (2006).

We proceed as follows. In Section 2, we introduce the formal preliminaries. We then define a general class of plausibility functions in Section 3, and we show how a sequence of plausibility functions can be used to represent an uncertain sequence of actions and observations. We refer to such a sequence as a *graded world view*, and we demonstrate how an agent’s beliefs about event histories can be captured by graded world views through taking aggregates of the constituent plausibility functions. In Section 4, we demonstrate that graded world views can be seen as epistemic states, and we use our basic framework to define belief change in domains involving uncertainty about the actions and observations that have occurred. We compare our approach with related work in Section 5, and we discuss limitations and advantages in Section 6. In Section 7, we offer some concluding remarks.

2. Preliminaries

In this section, we introduce some standard formal machinery for modelling belief change, and for reasoning about action effects. We also introduce a simple motivating example to illustrate the sort of problem we would like to address.

2.1 Belief Revision

Belief revision refers to the process in which an agent incorporates new information along with some pre-existing beliefs. The most influential approach to belief revision is the AGM approach (Alchourrón, Gärdenfors, & Makinson, 1985). Let \mathbf{F} denote a finite set of *fluent symbols*, which represent binary properties of the world that may change over time. For example, there may be a fluent symbol *Raining* $\in \mathbf{F}$ that is true just in case it is raining. A *state* is a propositional interpretation of \mathbf{F} , indicating which fluents are true and which are false. In the AGM approach to belief revision, the beliefs of an agent are represented by a *belief set*, which is a deductively closed set of formulas over \mathbf{F} . Since \mathbf{F} is finite, we can equivalently define the beliefs of an agent to be represented by a single formula.

Informally, a *belief revision operator* is a function that takes a belief set ϕ and a formula for revision γ as input and returns a new formula that represents a new belief set incorporating γ . An *AGM revision operator* is a binary function $*$ that satisfies the *AGM postulates*. The following reformulation of the AGM postulates is due to Katsuno and Mendelzon (1991). In the postulates, \equiv denotes logical equivalence.

- [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 captured semantically by introducing a formal notion of *plausibility*. For a particular agent, we say that a state of the world s is more plausible than another state s' if that agent is more likely to abandon belief in s' when presented with new information. It turns out that every AGM revision operator can be characterized by a class of plausibility orderings over the set of states. This has been proved for several different representations of plausibility, including total pre-orders over states (Katsuno & Mendelzon, 1991), systems of spheres (Grove, 1988), and ordinal conditional functions (Spohn, 1988).

A *belief state* is a set of states, informally the set of states that an agent considers possible. An *observation* α is also a set of states, which informally represents some peice of information that an agent receives that provides evidence that the actual state is in α . In this paper, we are primarily interested in belief change as a process on states rather than formulas; it is clear that the AGM approach can equivalently be defined on belief states.

While we have restricted the set of fluents and the set of actions to be finite, this is really just for convenience when discussing examples. Our formal model will be based on transition systems and ranking functions over sets. It would certainly be possible to define transition relations and ranking functions over infinite sets. Since our definition of an aggregate is very general, this would introduce no formal complications. However, from this point on, we will maintain the restriction to finite domains in order to simplify the discussion.

2.2 Transition Systems

We are interested in action domains that can be described by supplementing the set \mathbf{F} of fluent symbols with a finite set \mathbf{A} of action symbols and a transition system describing the effects of actions (Gelfond & Lifschitz, 1998). A pair (\mathbf{F}, \mathbf{A}) is called an *action signature*. While we have defined a *state* to be an interpretation over \mathbf{F} , it is often convenient to identify a state s with the *set* of fluent symbols true in s . We use this convention in the following definition, and throughout the rest of the paper.

Definition 1 *A transition system is a pair $\langle S, R \rangle$ where $S \subseteq 2^{\mathbf{F}}$, $R \subseteq S \times \mathbf{A} \times S$.*

We restrict attention to *deterministic* transition systems, i.e. we assume that $\langle s, A, s' \rangle \in R$ and $\langle s, A, s'' \rangle \in R$ implies $s' = s''$. We also assume that \mathbf{A} always contains a distinguished null action symbol denoted by λ .

Belief update is the belief change that occurs when an agent becomes aware of a change in the state of the world. Note that this is a distinct process from belief revision, which is typically understood to capture the belief change that occurs when an agent has obtained new information about an unchanged world. One highly influential approach to belief update is the Katsuno-Mendelzon approach (Katsuno & Mendelzon, 1992), which is superficially similar to AGM revision in that the new information to be incorporated is encoded as a propositional formula. By contrast, in this paper, we define update with respect to an action with effects given by an underlying transition system. In other words, we define belief update operators that take a belief state and an action as input, and return a new belief state that represents an agent's new beliefs after the action has been executed.

Definition 2 Let $T = \langle S, R \rangle$ be a transition system. The update function $\diamond : 2^S \times \mathbf{A} \rightarrow 2^S$ is given by $\alpha \diamond A = \{s \mid \langle s', A, s \rangle \in R \text{ for some } s' \in \alpha\}$.

We remark that this notion of belief change can also be called *belief progression* or *action progression*. Note that it is possible in the general case for the result of this update to be empty, in the case where the transition system does not include any outgoing edge labelled A from a state in α . In order to avoid this problem, it is often convenient to restrict attention to transition systems where every action has an outcome in every state. In practice, this can be achieved by assuming a self-loop for the action A at any state where no outcome state is given.

2.3 Belief Evolution

In previous work, we introduced so-called *belief evolution* operators to reason about alternating sequences of updates and revisions (Hunter & Delgrande, 2011). A belief evolution operator is defined with respect to a fixed AGM revision operator $*$ and a fixed update operator \diamond . For any action A and observation α , let $\alpha^{-1}(A)$ denote the set of all s such that $\{s\} \diamond A \subseteq \alpha$. Note that this set may be empty. The belief evolution operator \circ is defined such that, for any belief state κ ,

$$\kappa \circ \langle A, \alpha \rangle = (\kappa * \alpha^{-1}(A)) \diamond A.$$

The complete definition of \circ actually defines the belief change for a sequence of actions and observations; but the details of this definition are not required at present. The important point is that the way each observation is incorporated depends on the preceding actions.

The intuition behind belief evolution is that the final state must be a possible effect of the most recently executed action A . This intuition is satisfied in the original definition by revising the initial belief set prior to computing the effects of A . As a result, belief evolution operators have a non-Markovian character; an observation can not be incorporated just by considering the current state of the world. Instead, an agent incorporates a new observation by looking back at the original state together with the complete history of actions. However, this operation can also be understood in a Markovian manner if we allow the current belief state to include some representation of the states that can be “ruled out” based on the actions that have previously occurred. In this manner, we can define a Markovian form of belief revision that is equivalent to belief evolution (Hunter, 2014).

Belief evolution provides a reasonable model for problems like Moore’s litmus paper problem (Moore, 1985). In this problem, an agent dips a piece of litmus paper in a beaker to determine if the beaker contains an acid or a base. Hence, the agent is performing an action first and then observing the results. It seems, however, that the observed results should affect the agent’s *initial* belief state. That is, after the litmus paper turns red, the agent is likely to conclude that the beaker contained an acid even before dipping.

Belief evolution operators satisfy the following properties. In these properties, $A \in \mathbf{A}$, \bar{A} denotes a sequence of action symbols of indeterminate length, and $2^{\mathbf{F}}$ denotes the set of all states over \mathbf{F} . The symbols κ and α both range over $2^{\mathbf{F}}$, though we think of κ as a belief state and we think of α as an observation. Following the standard convention for such postulates, we implicitly quantify universally over all variables.

1. If $(2^F \diamond \bar{A}) \cap \alpha \neq \emptyset$, then $(\kappa \diamond \bar{A}) * \alpha \subseteq \alpha$
2. If $(2^F \diamond \bar{A}) \cap \alpha = \emptyset$, then $(\kappa \diamond \bar{A}) * \alpha = \kappa \diamond \bar{A}$
3. $(\kappa \diamond \bar{A}) \cap \alpha \subseteq (\kappa \diamond \bar{A}) * \alpha$
4. If $(\kappa \diamond \bar{A}) \cap \alpha \neq \emptyset$, then $(\kappa \diamond \bar{A}) * \alpha \subseteq (\kappa \diamond \bar{A}) \cap \alpha$
5. $(\kappa \diamond \bar{A}) * \alpha \subseteq 2^F \diamond \bar{A}$

Informally, these properties assert that an agent should respect the history of actions executed when incorporating a new observation. Note that we are assuming that $*$ is an AGM revision operator, which would appear to make properties (1) and (3) trivial. However, they are included as properties of belief evolution to emphasize the fact that the actions preceding the observation need to be considered in all cases where revision occurs (Hunter & Delgrande, 2011).

One of the main limitations of belief evolution is that it is not possible to represent erroneous action histories; it is assumed that the action history is always correct. This is a reasonable assumption in action domains where there is a single agent executing actions that can not fail. However, if exogenous or failed actions are permitted, then this assumption is difficult to support. In general, an agent may have incorrect beliefs about the actions that have occurred in the past. One of our aims in this paper is to generalize our previous work to allow erroneous action histories.

2.4 Motivating Example

We introduce a common-sense example in which an agent needs to compare the plausibility of certain actions with the plausibility of observations. We will return to this example periodically as we introduce our formal machinery.

Consider a simple action domain involving four agents: Bob, Alice, Eve, and Trent. Bob places a chocolate chip cookie on his desk and then leaves the room; he believes that no one is likely to eat his cookie while he is gone. At Time 1, Bob knows that Alice is at his desk. At Time 2, Bob knows that Eve is at his desk. After Eve leaves his desk, Trent comes and tells Bob that a bite has been taken from the cookie.

Given the preceding information, Bob can draw three reasonable conclusions: Alice bit the cookie, Eve bit the cookie, or Trent gave him poor information. If Bob has no additional information about the world, then each conclusion is equally plausible. However, we suppose that Bob does have some additional information. In particular, suppose that Alice is a close friend of Bob and they have shared cookies in the past. Moreover, suppose that Bob believes that Trent is always honest. Bob's additional information about Alice and Trent provides a sufficient basis for determining which of the three possible conclusions is the most plausible.

Informally, prior to Trent's report, Bob believes that his cookie was unbitten at all earlier points in time. After Trent tells him the cookie is bitten, he must determine the most plausible world history consistent with this information. In this case, the most plausible solution is to conclude that Alice bit the cookie. Note that this conclusion requires Bob to alter his subjective view of the action history. There is a non-monotonic character to

belief change in this context, because Bob may be forced to postulate and retract actions over time in response to new observations. The consequences of changing the action history are determined by the underlying transition system. In order to represent this kind of reasoning, we need to be able to compare the plausibility of action occurrences at different points in time.

This example illustrates the kind of action scenario that we would like to capture, because it requires an agent to determine the most plausible history given some a priori notion of plausibility for actions and observations. For instance, it is intuitively clear that this problem can not be resolved without comparing the plausibility of Trent’s report being accurate versus the plausibility of Alice biting the cookie. By addressing this issue in a straightforward manner, we demonstrate how existing actions formalisms can be extended and employed to handle similar situations.

There is possibly one contentious issue in our discussion of this example. The information provided by Trent is not really an “observation” in the usual sense of the word; instead, it is a report of information from an external source. However, in the sections that follow, we will treat reports and observations in the same manner. Specifically, we will capture both by a ranking function over states that indicates which states are supported by the observation/report and to what degree. Of course, in reality, observations and reports are quite different in that the manner in which a report is incorporated depends on *trust* in the reporting agent. The relationship between trust and belief change is a topic of current interest, which is generally handled by introducing some extra formal machinery to encode the trust held in another agent (Lorini, Jiang, & Perrussel, 2014; Hunter & Booth, 2015). It would certainly be possible to follow this approach in the present paper, using distinct formal tools to capture trust as a phenomenon that is distinct from the perceived accuracy of sensing. However, mathematically, we would like to end up with a single ranking over states nevertheless. As such, for the present paper, it is more convenient to represent reported information through a single “combined” ranking that captures the final plausibility attached to each state based on all considerations an agent might make.

3. Ranking Functions over Actions and States

The approach presented in this paper is based on a simple notion: when an agent is uncertain about some action or observation, resolving this uncertainty generally involves comparing the relative likelihood of possible alternatives. Our work is distinguished by the fact that we use *sequences of ranking functions* to represent uncertainty over both actions and states. We will see that sequences of ranking functions can capture many natural reasoning problems. By developing high-level representations of these problems, we can see that some notion of *magnitude* of likelihood is useful for reasoning about belief change with uncertain action histories. This allows us to compare our work with existing formalisms for reasoning about epistemic action effects, particularly those in which the representation of uncertainty is limited to orderings over states. Our fundamental goal is to demonstrate that, in some cases, sequences of quantitative plausibility orderings have a expressive advantage that is significant from the perspective of knowledge representation. The aim is to develop our approach at a high-level, in a manner that can easily be translated into other action formalisms.

3.1 Plausibility Functions

We are interested in action domains where an agent’s beliefs about an action history may be incorrect. In this context, the action that is believed to be executed at any given point in time can be represented by a total pre-order over possible actions. The minimal elements of such a pre-order represent the actions that were most likely executed, and moving higher in the ordering gives increasingly implausible possibilities. Representing actions in this manner allows an agent to determine plausible alternative actions in the face of conflicting evidence. Similarly, an agent needs a mechanism for ordering states in order to represent fallible observations and fallible beliefs. Moreover, we would like to be able to compare orderings over actions with orderings over states. One natural way to create mutually comparable orderings is by assigning quantitative plausibility values to every action and state at every point in time. Towards this end, we define plausibility functions.

Definition 3 *Let X be a non-empty set. A plausibility function over X is a function $r : X \rightarrow \mathbf{N}$.*

If r is a plausibility function and $r(x) \leq r(y)$, then we say that x is at least as plausible as y . We will only be interested in plausibility functions over finite sets, where there is always a non-empty set of minimally ranked elements.

Plausibility functions are inspired by Spohn’s ordinal conditional functions (Spohn, 1988), but there are some important differences. First, we allow plausibility functions over an arbitrary set X , rather than restricting attention to propositional interpretations. This allows us to treat actions in the same manner that we treat observations. Another important difference is that ordinal conditional functions must always assign rank 0 to a non-empty subset of elements of the domain. Plausibility functions are not restricted in this manner; the minimal rank for a given plausibility function may be greater than 0. We have defined plausibility functions in this manner because we will be interested in taking sums over plausibility functions, and we need to ensure that such sums also define plausibility functions.

We remark that Darwiche and Pearl also consider ranking functions that do not necessarily assign rank 0 to any states (Darwiche & Pearl, 1997). However, Darwiche and Pearl define the belief state associated with r to be the set of states that are assigned rank 0. Under this convention, ranking functions that never assign rank 0 are associated with the empty belief state. By contrast, we associate a non-empty belief state with every plausibility function.

We introduce some useful terminology and notation. Let r be a plausibility function over X . The minimum and maximum values obtained by r are denoted by \min_r and \max_r , respectively. We define $Bel(r)$ to be the set $\{s \mid r(s) = \min_r\}$. This notation is intended to suggest that $Bel(r)$ is the set of actions or states that are believed. To be clear, when we say that a state s is *believed*, we mean that as far as the agent is concerned, the actual world could be described by s . On the other hand, when we say that an action A is *believed*, we mean that the agent views A to be the action that was executed. Note that \min_r is always defined, because any non-empty set of natural numbers has a minimum element. On the other hand, \max_r is only guaranteed to be defined when X is finite. However, we will be only concerned with plausibility functions over states and actions; these are both finite sets according to our original definitions.

For $\alpha \subseteq X$, we define $r(\alpha)$ to be the minimum value obtained by r for some $s \in \alpha$. The *degree of strength* of a plausibility function r is the least n such that $\min_r + n = r(s)$ for some $s \notin Bel(r)$. Hence, the degree of strength of r is the span between the plausibility of the minimally ranked elements and the non-minimally ranked elements. If r has degree of strength n , this means that every $s \notin Bel(r)$ has a plausibility value that is at least n higher than $r(Bel(r))$. There are two natural interpretations of the degree of strength of a plausibility function r over a set of states. If we think of r as an initial epistemic state, then the degree of strength is an indication of how strongly it is believed that the actual state is in $Bel(r)$. If we think of r as an observation, then the degree of strength is a measure of the subjective reliability of r . In the case where X is a set of states, we use the terms “degree of strength” and “degree of belief” interchangeably.

The notion of degree of strength is of crucial importance in our approach. If we use a total pre-order to represent plausibility, then there is really no corresponding notion of strength to distinguish a situation where the most plausible belief is very strongly held versus one where it is weakly held. But this kind of distinction is essential when we have to sort out a sequence of events; we need some notion of strength of belief to decide which state or action is the hardest to give up. It is possible to represent this kind of information in an ordering with “empty levels,” of course. However, our focus will be on aggregating sequences of belief states and actions. In order to make this as simple as possible, we suggest that it is better to use the compact representation given by a quantitative ranking function that lends itself naturally to arithmetical combinations.

Note that Spohn (1988) defines the degree of strength of a subset of X , rather than the degree of strength of a ranking function. Our definition coincides with Spohn’s definition if we identify the degree of strength of r with Spohn’s degree of strength of the set $Bel(r)$. Hence, we use the same conception of degree of strength, but we are only interested in the strength of belief in the minimally ranked elements.

In order to illustrate the application of plausibility functions over different domains, we continue our simple example.

Example (cont’d) Let $\mathbf{F} = \{BiteTaken\}$ and let $\mathbf{A} = \{BiteAlice, BiteEve\}$. Both actions have the same effect, namely they both make the fluent *BiteTaken* become true. We represent the problem with 3 plausibility functions: a_1 , a_2 , and o_2 .

1. a_1 is a plausibility function over actions at Time 1
2. a_2 is a plausibility function over actions at Time 2
3. o_2 is a plausibility function over states at Time 2

Informally, each function should obtain a minimum value at the event that Bob considers the most plausible at the given point in time. Since Bob initially believes that no one will eat his cookie, both a_1 and a_2 should obtain a minimum value at the null action λ . Trent’s report that the cookie has been bitten at Time 2 is represented as a plausibility function over states, by defining o_2 with a minimum at the set of worlds where the cookie has a bite out of it. Note that we will generally treat reported information in this manner; the degree of strength of a report is an indication of trust in the agent providing the report. The

additional soft constraints about Bob’s relationships are used to determine the magnitude of the values for each event. Define a_1 and a_2 by the values in the following table.

	λ	<i>BiteAlice</i>	<i>BiteEve</i>
a_1	0	1	10
a_2	0	10	3

The columns of the table give the plausibilities of each action at each time. In particular, a_1 encodes that the plausibility λ occurs is 0, whereas the plausibility of *BiteAlice* and *BiteEve* are 1 and 10, respectively. The degree of strength of a_1 is 1, whereas the degree of strength of a_2 is 3. The fact that Alice is more likely to bite the cookie is represented by assigning a low plausibility value to *BiteAlice* at Time 1.

We define o_2 as follows, so give the plausibility values for all states at time 2.

	\emptyset	$\{BiteTaken\}$
o_2	9	0

Hence, the observation $\{BiteTaken\}$ is assigned the minimum plausibility value, and the only alternative observation is assigned a very high plausibility value. The degree of strength of o_2 is 9. This reflects the fact that Trent’s report is understood to be stronger than the assumption that Alice and Eve do not bite the cookie.

Note that the degree of strength of a_1 is less than the degree of strength of a_2 and o_2 . This gives an indication that Bob has comparatively less confidence in his beliefs about the action at Time 1.

Note that, in this example, we have explicitly referred to points in time. We do not, however, include any formal representation of time in the our framework; any reference to time should be interpreted as an informal explanatory device. We are concerned with reasoning about scenarios that involve sequences of actions where each action is instantaneous, and all actions are executed consecutively. Nevertheless, since we allow observations following null actions, we can imagine that actions are being executed in accordance with a bounded global clock that allows one action per tick.

3.2 Graded World Views

We define a *graded world view* to be an alternating sequence of plausibility functions over $2^{\mathbf{F}}$ and plausibility functions over \mathbf{A} . Hence, at each time i , we use a plausibility function over $2^{\mathbf{F}}$ to represent an agent’s beliefs about the state of the world and we use a plausibility function over \mathbf{A} to represent an agent’s beliefs about the action that occurs. Informally, a graded world view represents an agent’s subjective view of the evolution of the world in the context of imperfectly known action histories. The rationale behind using a sequence of ranking functions is to eventually make it possible to compare the likelihood of actions and observations at different points in the sequence, to arrive at the most plausible sequence of events. We have the following formal definition.

Definition 4 *A graded world view of length n is a $(2n + 1)$ -tuple*

$$\langle OBS_0, ACT_1, OBS_1, \dots, ACT_n, OBS_n \rangle$$

where each OBS_i is a plausibility function over $2^{\mathbf{F}}$ and each ACT_i is a plausibility function over \mathbf{A} .

At time i , the most plausible actions are the minimally ranked actions of ACT_i and the most plausible states are the minimally ranked states of OBS_i . We take OBS_0 to represent the initial belief state, and each subsequent OBS_i to represent a new observation. If $ACT = \langle ACT_1, \dots, ACT_n \rangle$ and $OBS = \langle OBS_0, \dots, OBS_n \rangle$, then we write $\langle ACT, OBS \rangle$ as a shorthand for the graded world view $\langle OBS_0, ACT_1, OBS_1, \dots, ACT_n, OBS_n \rangle$. Informally, a graded world view represents an agent's subjective view of the history of actions and observations.

We remark briefly on the intuition behind graded world views. We are interested in action domains involving actions that are both partially observable and fallible. However, for the moment we do not consider failed actions; we address this issue briefly in §4.6. The plausibility of an action A represents an agent's confidence that A was successfully executed at a given instant. Hence, the lowest plausibility values will be assigned to actions that an agent has executed, or actions that an agent has observed directly. Higher plausibility values will be assigned to exogenous actions that are assumed to be unlikely, or action occurrences that are only believed based on external reports. In §3.4, we provide some additional motivation for plausibility values by illustrating a correspondence with subjective probability functions.

Note that graded world views essentially represent the initial belief state as an observation at time 0. The underlying assumption in a graded world view is that the initial belief state is no different than any subsequent observation; there is no reason to automatically prefer the initial beliefs over new information, nor is there any reason to automatically disregard the initial beliefs given new information. In the definition of a graded world view, we have used indices in a manner that is not symmetric in order to emphasize the unique stature of OBS_0 . In particular, note that there is no ACT_0 in the definition. This simple notational convention is intended to highlight the fact that OBS_0 has a slightly different stature at an informal level.

3.3 Aggregate Plausibility Functions

Given a graded world view $\langle ACT, OBS \rangle$, we would like to be able to determine the most plausible history of the world. We formally define the notion of a *history* over a transition system.

Definition 5 Let $T = \langle S, R \rangle$ be a transition system. A history of length n is a tuple $\langle s_0, A_1, \dots, A_n, s_n \rangle$ where for each i :

1. $s_i \in S$,
2. $A_i \in \mathbf{A}$, and
3. $\langle s_i, A_{i+1}, s_{i+1} \rangle \in R$.

Let $HIST_n$ denote the set of histories of length n .

Note that $HIST_n \subseteq S \times (\mathbf{A} \times S)^n$. Ideally, we would like to use graded world views to assign plausibility values to histories. However, a graded world view does not provide

sufficient information to define a unique plausibility function over histories. For example, a graded world view does not indicate the relative weight of recent information versus initial information. In order to determine the most plausible history, we need some mechanism for combining a sequence of plausibility functions.

Although a graded world view does not define a unique plausibility function over histories, we can define a general notion of consistency between graded world views and plausibility functions on histories. Let r_0, \dots, r_n be plausibility functions over X_0, \dots, X_n , respectively. Let r be a plausibility function over $X_0 \times \dots \times X_n$. We say that r is *consistent with* $\langle r_0, \dots, r_n \rangle$ if, for every i and every $x_i, x'_i \in X_i$

$$\begin{aligned} r_i(x_i) < r_i(x'_i) \\ \iff \\ r(\langle x_0, \dots, x_i, \dots, x_n \rangle) < r(\langle x_0, \dots, x'_i, \dots, x_n \rangle) \end{aligned}$$

So r is consistent with $\langle ACT, OBS \rangle$ just in case r increases monotonically with respect to each component of $\langle ACT, OBS \rangle$. Any plausibility function r that is consistent with $\langle ACT, OBS \rangle$ provides a potential candidate ranking over histories.

Define an *aggregate plausibility function* to be a function that maps every graded world view of length n to a plausibility function over $HIST_n$. We are interested in aggregate plausibility functions in which the output is always consistent with the input. Hence, we say that an aggregate plausibility function agg is *admissible* if, for every $\langle ACT, OBS \rangle$, the function $agg(\langle ACT, OBS \rangle)$ is consistent with $\langle ACT, OBS \rangle$.

We provide some examples. Note that aggregate plausibility functions return a function as a value; we can specify the behaviour of an aggregate by specifying a plausibility value for each pair consisting of a graded world view and a history. Let $h = \langle s_0, A_1, \dots, A_n, s_n \rangle$. One admissible aggregate is obtained by taking the sum of plausibility values.

$$sum(\langle ACT, OBS \rangle)(h) = \sum_{i=1}^n ACT_i(A_i) + \sum_{i=0}^n OBS_i(s_i)$$

A weighted sum can be used to reflect the relative importance of different time points. For each i , let b_i be a positive integer.

$$sum_s(\langle ACT, OBS \rangle)(h) = \sum_{i=1}^n ACT_i(A_i) + \sum_{i=0}^n b_i \cdot OBS_i(s_i).$$

By setting $b_i = 2^i$, the aggregate function sum_s can be used to represent a strict preference for recent observations. This is a standard assumption in many approaches to belief revision. We could add a similar weight to the action histories as well, which would give another distinct aggregate. The functions sum and sum_s are just two simple examples; many more examples can be defined by specifying aggregate functions that increase monotonically with each component.

We return to the cookie example to illustrate how the reasoning involved can be captured with graded world views and aggregate plausibility functions.

Example (cont'd) We have already defined plausibility functions a_1, a_2 and o_2 . In order to give a complete graded world view, we need to define two more plausibility functions over

states. In particular, we need to give a plausibility function o_0 representing Bob’s initial beliefs and we need to give a plausibility function o_1 representing the null observation that Bob makes at Time 1.

First, we reiterate the description of a_1 and a_2 in the following table.

	λ	<i>BiteAlice</i>	<i>BiteEve</i>
a_1	0	1	10
a_2	0	10	3

The fact that Alice is more likely to bite the cookie is represented by assigning a lower plausibility value to *BiteAlice* at Time 1.

The plausibility function o_0 should assign a minimum value to the state where the cookie is unbitten. The plausibility function o_1 should assign the same value to every state. The plausibility function o_2 (given previously) represents Trent’s report that the cookie has been bitten. As noted previously, we treat reported information as an observation, and we use the degree of strength of the reported information as an indication of the reliability of the source. In this case, the degree of strength of o_2 is an indication of trust in Trent. We define o_0, o_1, o_2 in the next table.

	\emptyset	$\{BiteTaken\}$
o_0	0	9
o_1	0	0
o_2	9	0

Note that the degree of strength of o_2 is higher than the degree of strength of a_1 or a_2 . This reflects the fact that Trent’s report is understood to supersede the assumption that Alice and Eve do not bite the cookie. Graded world views have been defined precisely for this kind of comparison between action plausibilities and state plausibilities.

If we use the aggregate function *sum*, then we are interested in finding the minimal sum of plausibilities over $\langle o_0, a_1, o_1, a_2, o_2 \rangle$. By inspection, we find that the minimum plausibility is obtained by the following history:

$$h = \langle \emptyset, BiteAlice, BiteTaken, \lambda, BiteTaken \rangle.$$

This history represents the sequence of events in which Alice bites the cookie at time 1. Intuitively, this is the correct solution: given the choice between Alice and Eve, Bob believes that Alice is the more plausible culprit.

We remark that graded world views bear a resemblance to the generalized belief change framework proposed by Liberatore and Schaerf (2000). However, the Liberatore-Schaerf approach associates a “penalty” with state change, which is minimized when determining plausible models. As such, it is difficult to represent problems where non-null actions are strictly more plausible than null actions. By contrast, graded world views have no implicit preference for null actions. Moreover, our approach differs in that we allow actions with conditional effects given by a transition system.

3.4 Subjective Probabilities

One issue that arises from our definition of a graded world view is the fact that it is not clear how plausibility values should be assigned in practical problems. We address this problem by illustrating a correspondence between plausibility functions and *probability functions*. We simplify the discussion by restricting attention to rational-valued probability functions as follows.

Definition 6 *Let X be a non-empty set. A probability function over X is a function $Pr : X \rightarrow \mathbf{Q}$ such that*

- for all $x \in X$, $0 \leq Pr(x) \leq 1$
- $\sum_{x \in X} Pr(x) = 1$.

We do not need any other axioms of probability theory for our present purposes. At a common-sense level, it is clear what it means to say that “action A occurred at time t with probability p .” By contrast, the problem with plausibility values is that there is no obvious sense of scale; it is difficult to assign numerical plausibility values, because the numbers have no clear meaning. We illustrate how probability functions can be translated uniformly into plausibility functions, thereby giving a sense of scale and meaning to plausibility values.

Let Pr be a probability function over a finite set X . Let Q denote the least common denominator of all rational numbers $\frac{p}{q}$ such that $Pr(x) = \frac{p}{q}$ for some $x \in X$. Define the plausibility function r as follows.

1. If $Pr(x)$ is minimal, set $r(x) = Q$.
2. Otherwise, if $Pr(x) = \frac{p}{Q}$, then set $r(x) = Q - p$.

Hence, every probability function can be translated into a plausibility function.

Example (cont’d) Consider the following probability functions for the cookie example.

	λ	<i>BiteAlice</i>	<i>BiteEve</i>
Pr_{a_1}	.5	.45	.05
Pr_{a_2}	.5	.15	.35

	\emptyset	$\{BiteTaken\}$
Pr_{o_0}	.9	.1
Pr_{o_1}	.5	.5
Pr_{o_2}	.1	.9

The corresponding plausibility functions are given in the following tables.

	λ	<i>BiteAlice</i>	<i>BiteEve</i>
a'_1	10	11	20
a'_2	10	20	13

	\emptyset	$\{BiteTaken\}$
o'_0	1	10
o'_1	2	2
o'_2	10	1

It is easy to see that these plausibility functions are obtained from the plausibility functions given earlier by adding a constant to each value. In §4.2, we illustrate that adding a constant in this manner does not affect the class of minimally ranked histories.

Connecting plausibility functions with subjective probabilities provides some justification for the use of the aggregate function *sum*. In particular, if we assume that the subjective

probability functions are *independent*, then the probability of a given sequence of events is determined by taking a product. In the cookie example, we can compare the probability of Alice biting the cookie versus Eve biting the cookie:

1. $Pr(\langle \emptyset, BiteAlice, BiteTaken, \lambda, BiteTaken \rangle)$
 $= .9 \times .45 \times .5 \times .5 \times .9 = .091125$
2. $Pr(\langle \emptyset, \lambda, \emptyset, BiteEve, BiteTaken \rangle)$
 $= .9 \times .5 \times .5 \times .35 \times .9 = .070875$

It is easy to check that the history where Alice bites the cookie is actually the most probable history. So, in this example, the minimally ranked history according to the aggregate function *sum* is also the most probable history according to the sequence of probability functions. This is a general property of our translation: maximizing probability over independent probability functions corresponds to minimizing the sum over plausibility values. This follows simply from the fact that the more probable events have been assigned minimal plausibility values, and the fact that summation over natural numbers is an increasing function, whereas multiplication on fractions less than 1 is decreasing.

The correspondence described in this section relies on the assumption that sequences of actions are *independent*. It is worth noting, however, that this very often is not the case in practice. Instead, it is often the case that actions occur in sequences with some clear dependence between the individual actions. However, this is not our concern at present. The only reason for considering probability functions is to provide some intuition or motivation for the way that plausibility values can be assigned.

3.5 The Summation Convention

In order to ground the discussion, it is useful to choose a fixed aggregate function for assigning plausibility values to histories. As such, unless otherwise indicated, we will assume that plausibility values are assigned to histories by the aggregate function *sum*. Although this is not the only approach to combining plausibility functions, it provides a simple admissible aggregate function that is appropriate in many cases. In particular, we saw in the previous section that *sum* is appropriate for domains where the plausibility functions have been obtained from independent subjective probabilities.

We introduce some notation that will simplify the results in the next few sections. Recall that $sum(\langle ACT, OBS \rangle)$ is a plausibility function on histories. When the underlying graded world view is clear from the context, we will write $plaus(h)$ as a shorthand for $sum(\langle ACT, OBS \rangle)(h)$.

It is useful to introduce an operator that maps a graded world view to the most plausible histories.

Definition 7 Let WV_n denote the set of graded world views of length n for a fixed action signature. Define $\Phi : WV_n \rightarrow 2^{HIST_n}$ as follows:

$$\Phi(\langle ACT, OBS \rangle) = \{h \mid plaus(h) \leq plaus(g) \text{ for all } g \in HIST_n\}.$$

We have the following obvious equivalence

$$\Phi(\langle ACT, OBS \rangle) = Bel(sum(\langle ACT, OBS \rangle)).$$

It is also useful to use *plaus* to define a plausibility function over states.

Definition 8 *Let $\langle ACT, OBS \rangle$ be a graded world view. For any state s , define*

$$plaus_state(\langle ACT, OBS \rangle)(s)$$

to be the least n such that $plaus(\langle ACT, OBS \rangle)(h) = n$ for some history h with final state s .

So the plausibility of the state s is the rank of the most plausible history ending with s . When the underlying graded world view is clear from context, we simply write $plaus_state(s)$ for the plausibility of the state s . We extend the operator $Bel(\cdot)$ to graded world views by defining $Bel(\langle ACT, OBS \rangle)$ to be $Bel(plaus_state(\langle ACT, OBS \rangle))$. Hence, Bel takes a graded world view as an argument and returns the most plausible set of terminal states.

3.6 A Uniform Representation

Before considering applications and formal results, we briefly discuss the most novel and significant features of our approach. The notion of plausibility and the relation with theory change has been explored extensively in the literature. Work on nonmonotonic consequence operators, for example, has been informed by notions of preferred models (Kraus, Lehmann, & Magidor, 1990) and conditional knowledge (Lehmann & Magidor, 1992). There has also been a great deal of research on different representations of plausibility, ranging from orderings (Baltag & Smets, 2006, 2008; Britz & Varzinczak, 2013), to quantitative approaches to possibility (Benferhat, Dubois, & Prade, 1999), to variations on probabilistic models (Friedman & Halpern, 2001). We take a moment to position our work in this context.

As we have stated several times, what sets our work apart is that we apply the same notion of plausibility uniformly to both states and actions. This contrasts with most existing work, where a model of plausibility is used for representing initial beliefs, but then actions are handled with different formal machinery. In many cases, actions are atomic; belief change due to action is concerned with how a complicated plausibility structure changes in response to a distinct action occurrence.

In our model, an action occurrence is just another plausibility function. This has several advantages, which we show in the next sections. First, it allows us to easily compare strength of belief in a state with strength of belief in an action occurrence. For example, if I believe the lamp is on but I also believe that I toggled the switch: what should I believe? Situations of this form are common, and they require a comparison between two different forms of likelihood. In general, there is no natural logical apparatus that can help resolve this problem; we need some extra-logical information about which belief is stronger, the belief in the state or the belief in the action. There is a natural intuition that we actually need to be given some rankings to make this determination; that is exactly what we do here.

The second advantage of using a uniform representation of plausibility for actions is that it allows us to consider alternative actions. Looking at the literature on belief change and preferential models, it is clear that we actually need orderings over states in order to perform many kinds of reasoning. The reason we need such orderings is because we need to be able to specify the “best” alternatives to the things we believed to be true. We argue that the same is true for actions. When we find out that an action did not occur, we need to

have some mechanism for determining the next “best” alternative. We will see in the next section that there are many practical examples where this is important, and these examples can all be captured in a straightforward way by plausibility functions.

While it is tempting to look at our model and try to position it in the context of alternative models of plausibility, we argue that this is not the best way to look at this work. In fact, we could propose an approach very similar to the one in this paper based on sequences of probability functions with suitable aggregates. The fact that we are using Spohn-style ranking functions in particular is not the most important point. What is important is that we are using a measure of plausibility uniformly for observations and actions, and the plausibility measure has the feature that *differences* in plausibility have a magnitude. This allows us to determine which alternatives to our current beliefs about states and actions should be abandoned, because we can appeal to our single notion of magnitude to determine some notion of *minimal change*.

4. Using Graded World Views

In this section, we consider the most basic approach that one can use to try to find a minimally ranked history.

4.1 Pointwise Minima

Let $W = \langle ACT, OBS \rangle$ be a graded world view where

$$ACT = \langle ACT_1, \dots, ACT_n \rangle$$

and

$$OBS = \langle OBS_0, \dots, OBS_n \rangle.$$

The easiest way to determine a minimally ranked history is to simply take the most plausible actions and the most plausible worlds at each point in time. The following definition makes this notion more precise.

Definition 9 *Given a history $h = \langle s_0, A_1, \dots, A_n, s_n \rangle$, we say h is a pointwise minimum for $\langle ACT, OBS \rangle$ if, for all i ,*

1. *for all $A \in \mathbf{A}$, $ACT_i(A_i) \leq ACT_i(A)$, and*
2. *for all $s \in 2^{\mathbf{F}}$, $OBS_i(s_i) \leq OBS_i(s)$.*

The following proposition states that, if a graded world view has any pointwise minima, then those will be the most plausible histories.

Proposition 1 *Let $W = \langle ACT, OBS \rangle$ be a graded world view and let M be the set of pointwise minima for W . If $M \neq \emptyset$, then $\Phi(W) = M$.*

Proof It is sufficient to note that, for $h \in M$, $plaus(h) \leq plaus(g)$ for all histories g . \square Note, however, that histories are restricted in that each world must be the outcome of the preceding action. As such, it is possible that a graded world view will have no pointwise minimum. This is why the preceding proposition starts with the assumption that the set

of pointwise minima is non-empty. In cases where there are no pointwise minima, there will still be minimally plausible histories.

Finding pointwise minima is not easy in the general case.

Proposition 2 *Determining if a graded world view has a pointwise minimum is NP-complete.*

Proof Given a history, it is clear that checking if it is a pointwise minimum can be done in linear time; so the problem lies in *NP*. Assume a fixed set \mathbf{F} of fluent symbols. Let $W = \langle ACT, OBS \rangle$ where OBS has length $2^{|\mathbf{F}|}$, and assume that each OBS_i obtains a minimum at a distinct interpretation of \mathbf{F} . Then a pointwise minimum for W will correspond to a Hamiltonian path for the underlying transition system. The result follows since finding a Hamiltonian path is NP-complete. \square

The fact that it is already intractable to find a pointwise minimum suggests that finding minimal histories over a complicated aggregate function is likely to be computationally difficult. This is not a major concern for our present purposes however, as we view graded world views as a high-level tool to capture a variety of belief change scenarios. If one is interested in modelling concrete action domains, then it is important to choose aggregate functions that are well understood and computationally easy to minimize.

4.2 Equivalence

Clearly it is possible for two distinct graded world views to have the same set of minimally ranked world histories. In fact, it is possible for two distinct graded world views to induce the same preference ordering over histories. In this section, we define a natural equivalence relation over graded world views with an eye towards categorical representations. We start by defining a relation on plausibility functions.

Definition 10 *Let r_1 and r_2 be plausibility functions over a set X . We say that $r_1 \cong r_2$ if, for every $x, y \in X$,*

$$r_1(x) - r_1(y) = r_2(x) - r_2(y).$$

It is easy to verify the following: \cong is an equivalence relation, r_1 and r_2 have the same degree of strength, and $r_1 \cong r_2$ implies that $Bel(r_1) = Bel(r_2)$. Essentially, the only difference between equivalent plausibility functions is the minimum value; we make this idea more precise by defining a notion of normalization for plausibility functions.

Let r be a plausibility function. For any integer $z \geq -\min_r$, the *translation* of r by z is the plausibility function $x \mapsto r(x) + z$. It is easy to prove that $r \cong r'$ if and only if r' is a translation of r . We define the *normalization* of r to be the translation by $-\min_r$. The normalization of r is the unique plausibility function equivalent to r that obtains a minimum of 0.

We can extend the notion of equivalence to graded world views.

Definition 11 *Let W_1 and W_2 be graded world views over histories for a fixed action signature with a given underlying transition system. We say that $W_1 \cong W_2$ if, for every pair of histories g and h ,*

$$sum(W_1)(g) - sum(W_1)(h) = sum(W_2)(g) - sum(W_2)(h).$$

Unlike plausibility functions, it is possible to construct equivalent pairs of graded world views that are not obtained by translations.

The following proposition illustrates that every graded world view is equivalent to a graded world view consisting of normalized plausibility functions.

Proposition 3 *Let $\langle ACT, OBS \rangle$ be a graded world view. If $\langle ACT', OBS' \rangle$ is obtained by normalizing each component of ACT and OBS , then*

$$\langle ACT, OBS \rangle \cong \langle ACT', OBS' \rangle.$$

Proof Let g, h be histories. For ease of readability, let $plaus_1$ and $plaus_2$ denote $sum(\langle ACT, OBS \rangle)$ and $sum(\langle ACT', OBS' \rangle)$, respectively. Hence, $plaus_1$ and $plaus_2$ are functions over histories, obtained by minimization of the total sum. As such, each of $plaus_1$ and $plaus_2$ must have a *minimum* value, corresponding to the lowest possible sum over terms. Then the following equalities are immediate:

$$\begin{aligned} plaus_2(g) - plaus_2(h) &= plaus_1(g) - \min_{plaus_1} -plaus_1(h) + \min_{plaus_1} \\ &= plaus_1(g) - plaus_1(h). \end{aligned}$$

□

Hence, although we allow plausibility functions with minimum values larger than 0 in a graded world view, we can always pass to an equivalent graded world view consisting of normalized plausibility functions. We remark, however, that a graded world view defined by a sequence of normalized plausibility functions need not obtain a minimum of 0. In this case, the minimum will be 0 if and only if the graded world view has a pointwise minimum. It is also important to note that Proposition 3 only holds under the aggregate plausibility function *sum*.

4.3 Representing Belief States

Plausibility functions can be defined that simply pick out a distinguished set of elements of the domain. If $\alpha \subseteq X$ and c is an integer, let $\alpha \uparrow c$ denote the function from X to the set of integers that is defined as follows:

$$\alpha \uparrow c (s) = \begin{cases} 0 & \text{if } s \in \alpha \\ c & \text{otherwise} \end{cases}$$

If c is a positive integer, then $\alpha \uparrow c$ denotes a plausibility function in which the elements of α are the most plausible, and everything else is equally implausible. Plausibility functions of the form $\alpha \uparrow c$ will be called *simple*. If X is a set of states, then simple plausibility functions correspond to belief states; if X is a set of actions, then simple plausibility functions pick out the actions that are believed to have occurred. Using the terminology introduced earlier, we say that α is held with degree of belief c .

If $c > 0$, then $\alpha \uparrow -c$ does not actually define a plausibility function. However, allowing negative values leads to a simple symmetry in our notation. In the following proposition, $\tilde{\alpha}$ denotes the complement of α under set difference. So if α is a set of states, then $\tilde{\alpha} = 2^F - \alpha$.

Proposition 4 For any set α and positive integer c

$$\alpha \uparrow c \cong \tilde{\alpha} \uparrow -c.$$

Proof Let s, t be states. By definition, we have

$$\alpha \uparrow c (s) - \alpha \uparrow c (t) = \begin{cases} c & \text{if } s \in \alpha, t \notin \alpha \\ -c & \text{if } s \notin \alpha, t \in \alpha \\ 0 & \text{otherwise.} \end{cases}$$

and

$$\tilde{\alpha} \uparrow -c (s) - \tilde{\alpha} \uparrow -c (t) = \begin{cases} -(-c) & \text{if } s \notin \tilde{\alpha}, t \in \tilde{\alpha} \\ -c & \text{if } s \in \tilde{\alpha}, t \notin \tilde{\alpha} \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, the right hand sides of each equality are the same. \square

Suppose that

$$ACT = \langle ACT_1, \dots, ACT_n \rangle$$

and

$$OBS = \langle OBS_0, \dots, OBS_n \rangle$$

where each ACT_i and OBS_i is simple, with the same maximum plausibility c . This means that each ACT_i and OBS_i assigns either the value 1 or the value c to every element in their respective domains. In this case, we essentially have belief states with no plausibility ordering. In this case, it is easy to show that

$$\langle s_0, A_1, \dots, A_n, s_n \rangle \in \Phi(\langle ACT, OBS \rangle)$$

if and only if the cardinality of

$$\{A_i \mid A_i \in ACT_i\} \cup \{s_i \mid s_i \in OBS_i\}$$

is maximal among all histories. In other words, the most plausible histories are those that agree with $\langle ACT, OBS \rangle$ at the highest number of components. This is a reasonable approach to take in the trivial case where we have no prior ranking over states or actions. We emphasize that this is just a special case when the plausibility functions are simple and they share the same degree of strength. In the next section, we define belief change operations in terms of a general concatenation on plausibility functions. If we do not restrict the plausibility functions as we have here, then we get a much wider range of possible belief change operations.

4.4 Graded World Views as Epistemic States

An *epistemic state* is a representation of an agent's belief state that defines a total pre-order \preceq over all states (Darwiche & Pearl, 1997). If $s \preceq t$, then the underlying agent believes that it is more likely that the actual state of the world is s than t . The current belief state is given by the set of \preceq -minimal states. Recall also that a graded world view defines

a plausibility function *plaus.state* over states. So a graded world view clearly defines an ordering over states, and we can think of a graded world view as defining an epistemic state. The worlds that receive minimal rank in a graded world view are the worlds that are supported by the most reliable observations and actions. Using this ranking to define a plausibility ordering is tantamount to assuming that the plausibility of *s* is completely determined by the reliability of the source reporting that *s* occurs.

By viewing graded world views as epistemic states, we can define belief change operations in a more familiar manner. In particular, we can define belief change through a simple concatenation operator \bullet on graded world views. Given a sequence of plausibility functions $\bar{r} = \langle r_1, \dots, r_n \rangle$ and a plausibility function *r*, we let $\bar{r} \cdot r$ denote the sequence $\langle r_1, \dots, r_n, r \rangle$. Let $\langle ACT, OBS \rangle$ be a graded world view, let r_A be a plausibility function over actions and let r_S be a plausibility function over states. Define \bullet as follows:

$$\langle ACT, OBS \rangle \bullet \langle r_A, r_S \rangle = \langle ACT \cdot r_A, OBS \cdot r_S \rangle.$$

In this context the initial epistemic state is given by $\langle ACT, OBS \rangle$, which represents an agent's a priori beliefs about the history of observed actions and states. New actions and observations are incorporated by simply concatenating the new plausibility functions on to the initial graded world view. The new graded world view can be used to define a new ordering on histories through some aggregate function. For example, if we are using *sum* as the default aggregate, then the new ordering is immediate. Note that the new graded world view obtained in this manner also includes all historical information required for future belief change. As a special case of this simple concatenation operation, we get a new approach to update. For any set *X*, let 0 denote the plausibility function that uniformly assigns 0 to every element of *X*. We can identify the update $\langle ACT, OBS \rangle \diamond r_A$ with the following operation:

$$\langle ACT, OBS \rangle \bullet \langle r_A, 0 \rangle.$$

We can also define a natural approach to revision in this manner. Let *null* denote a plausibility function that assigns plausibility 0 to the null action λ , and assigns everything else a plausibility larger than the maximum value obtained by *sum*($\langle ACT, OBS \rangle$). We identify the revision $\langle ACT, OBS \rangle * r_S$ with the following operation:

$$\langle ACT, OBS \rangle \bullet \langle null, r_S \rangle.$$

Using plausibility functions to represent observations allows us to represent some natural problem domains that can not be easily represented if we restrict observations to sets of possible worlds. In particular, consider an action domain in which observations have varying degrees of reliability. In such domains, when an agent makes an observation that is inconsistent with the current belief state, there are two factors that should be considered: the strength of belief in the current belief state and the reliability of the observation. There is an obvious conflict that arises if we attempt to address both factors simultaneously. For example, suppose that the underlying agent strongly believes that *s* is a possible state of the world. Now suppose that the agent makes two observations.

1. One observation suggests that *s* is possible, but comes from an unreliable source.

2. Another observation suggests that s is not possible, and it comes from a very reliable source.

It can be difficult to determine the appropriate belief change in this scenario, particularly if strength of belief and observational reliability are treated independently. By quantifying the reliability of every observation, graded world views make it easy to resolve this kind of issue. We remark that problems of this form have also been addressed through the use of prioritized merging operators (Delgrande, Dubois, & Lang, 2006).

Note that there is an asymmetry in our definition of revision and update through the \bullet operator. In the case of update, we assume that the final observation assigns the same plausibility to every state. The symmetric definition for a single observation would be defined as follows:

$$\langle ACT, OBS \rangle \bullet \langle 0, r_S \rangle.$$

However, this definition allows an arbitrary action to occur immediately before the observation. If we want to assume that the graded world view $\langle ACT, OBS \rangle$ gives a complete picture of the world at the time of the observation, then we need to assume that any intermediary action is null. Hence, the asymmetry is not due to any significant difference between actions and observations; the asymmetry is simply due to the fact that graded world views involve alternating sequences of actions and observations, with actions occurring first by default.

In this section we have illustrated that a graded world view defines an epistemic state. If we take an epistemic state to be a total pre-order on states, then the converse is clearly false: an ordering on states does not provide enough information to define numerical ranking functions over states. The move from epistemic states to graded world views is motivated by the same kind of concern that motivates the move from belief states to epistemic states. In particular, belief states in AGM revision can be understood to represent the minimal elements in some ordering of states. Hence, a belief state can provide a partial description of an ordering, and an ordering can in turn provide a partial description of a graded world view. A belief state is sufficient for single-shot revision, provided that an ordering is implicit in the revision operator. However, a belief state is not sufficient if we need to explicitly reason about the way plausibility orderings are modified. Similarly, orderings on states are sufficient for reasoning about preferences over states, but they are not sufficient if we need to explicitly reason about action histories. In the next section, we clarify this point through practical examples.

4.5 Representing Natural Action Domains

We illustrate how some interesting phenomena can be represented by graded world views. The simplest examples involve graded world views of length 1. In particular, we initially focus on graded world views of the form

$$\langle INIT \rangle \bullet \langle r_A, r_S \rangle.$$

In this context, $INIT$ represents the initial belief state of an agent, r_A represents an agent's beliefs about the action that has been executed, and r_S represents the observed state of the world. Previously we have used OBS_0 for the initial plausibility function over states; the change in notation is just to emphasize that $INIT$ is some a priori initial ordering

over states. To be clear, $INIT$, r_A , and r_S are all *plausibility functions*. As such, we can define the degree of strength of each. To facilitate the exposition, we denote the degrees of strength by $deg(INIT)$, $deg(r_A)$, and $deg(r_S)$ respectively. Varying the magnitudes of these values allows us to capture several different underlying assumptions.

1. Fallible initial beliefs: $deg(INIT) < deg(r_A)$ and $deg(INIT) < deg(r_S)$.
2. Erroneous perception: $deg(r_S) < deg(INIT)$ and $deg(r_S) < deg(r_A)$.
3. Fallible action history: $deg(r_A) < deg(INIT)$ and $deg(r_A) < deg(r_S)$.

As a simple example, suppose that an agent believes a certain lamp is initially on, then the power switch is toggled, and then the agent observes that the lamp is actually still on. Clearly this sequence of events can not consistently be believed by a rational agent. Manipulating the degrees of strength of $INIT$, r_A and r_S gives an agent some mechanism for resolving such conflicts. In case (1), the agent is not completely certain that the lamp was initially on. As such, the easiest way to incorporate the new information is to change the initial belief state. By contrast, in case (2), the agent is not completely certain that the lamp is still on after toggling the switch. In this case, since the agent is confident the lamp was initially on and the switch was toggled, it is natural to reject the observation and believe that the lamp is now off. The distinction between these two cases cannot be captured without some notion of reliability. In case (3), the agent would resolve the conflict by believing that the attempt to toggle the power switch had failed.

The special case in which the degree of strength is 0 also captures some important phenomena. Note that a plausibility function r has degree of strength 0 just in case there is some constant c such that $r(x) = c$ for all x . As such, a degree of 0 indicates that every element of the domain receives minimal rank. We consider the informal interpretation of a degree 0 for each plausibility function in our schematic example.

1. If $deg(INIT) = 0$, then every initial state is equally plausible. The agent has no contingent a priori beliefs about the state of the world.
2. If $deg(r_O) = 0$, then r_O represents a null observation. The observation OBS does not provide evidence for any particular state.
3. If $deg(r_A) = 0$, then every action is equally likely. So the agent is completely ignorant about the action that has occurred, and we can think of r_A as an exogenous action beyond the agent's control.

These are relatively crude distinctions, but they still capture important classes of problems. Roughly speaking, the problems that we have addressed thus far can be captured by a plausibility ordering over sequences of the form

$$\kappa \diamond A_1 * \alpha_1 \diamond \dots \diamond A_n * \alpha_n$$

where κ is a belief state, each A_i is an action symbol, and each α_i is an observation.

For the purpose of comparison, we remark that belief evolution operators defined in Section 2.3 are only useful for problems in which the underlying plausibility ordering is

given as follows, for some permutation p_1, \dots, p_n of $1, \dots, n$.

$$\left. \begin{array}{c} A_1 \\ \vdots \\ A_n \end{array} \right\} \prec \alpha_{p_1} \prec \alpha_{p_2} \prec \dots \prec \alpha_{p_n}$$

By contrast, graded world views are suitable for any total pre-order over $A_1, \alpha_1, \dots, A_n, \alpha_n$. But this is not the entire class of problems representable by graded world views. By using a ranking function for each event, we are able to draw two additional distinctions that can not be represented by a simple ordering. First, we are able to represent changes in plausibility that do not affect the ordering of states. This is useful for representing action domains where an agent must observe a single piece of evidence multiple times before believing it is correct. Second, we are able to represent *graded evidence*; we use the term graded evidence to describe situations where an observation actually supports several different conclusions with different degrees of confidence. We conclude this section with two examples illustrating action domains that are difficult to represent if we only have an ordering over the plausibility of events.

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.

This example can easily be represented by a graded world view as follows. We assume that the underlying action signature contains, among others, a fluent symbol *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 define a graded world view that represents this action domain.

1. $OBS_0 = ON \uparrow 10$
2. $ACT_1 = \{TurnLampOff\} \uparrow 3$
3. $OBS_1 = ON \uparrow 2$
4. $ACT_2 = \lambda \uparrow 10$
5. $OBS_2 = ON \uparrow 2$

Note that $\Phi(\langle OBS_0, ACT_1, OBS_1 \rangle)$ consists of all histories where the lamp is turned off at time 1. However, $\Phi(\langle OBS_0, ACT_1, OBS_1, ACT_2, OBS_2 \rangle)$ consists of all histories where the lamp is not turned off at time 1. 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 approximately \$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 to be believable, 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 actually have been \$8.

To represent this example, we first define $ACT_1 = \lambda \uparrow 10$ because Bob believes that no ontic actions have occurred. We assume that there are fluent symbols $Cost1, Cost2, \dots, Cost9$ interpreted to represent the cost of the gift. We define a plausibility function OBS_0 representing Bob’s initial beliefs.

$$OBS_0(s) = \begin{cases} 0 & \text{if } s = \{Cost7\} \\ 1 & \text{if } s = \{Cost6\} \text{ or } s = \{Cost8\} \\ 3 & \text{otherwise} \end{cases}$$

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 OBS_1 representing the observation of the receipt.

$$OBS_1(s) = \begin{cases} 0 & \text{if } s = \{Cost3\} \\ 1 & \text{if } s = \{Cost8\} \\ 3 & \text{otherwise} \end{cases}$$

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 state of the world is the state in which the price is \$8. In order to draw this conclusion, Bob needs observations that provide 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 natural common-sense reasoning problems in which an agent needs to consider aggregate plausibilities over a sequence of actions and observations. Graded world views are well-suited for reasoning about such problems.

4.6 Non-deterministic and Failed Actions

In this section, we consider actions with non-deterministic effects. We remark, in particular, that fallible actions can be understood as actions with non-deterministic effects; fallible actions are just actions where some possible outcomes are considered to be failures. In the simplest case, for example, a failed action might be one which leaves the state of the world unchanged. Hence, by addressing non-deterministic effects, we are also handling actions that might fail. Our basic approach is the following. We introduce some new machinery for the representation of non-deterministic actions, and then we demonstrate that the new machinery is unnecessary when we use summation to determine the plausibility of histories. As such, we can reasonably restrict attention to deterministic actions when proving formal expressibility results for graded world views.

Given a non-deterministic transition system $T = \langle S, R \rangle$ and a graded world view W , it is not clear how we should choose the effects of each action in the most plausible world histories. This problem can be solved by attaching a plausibility value to the possible effects of each action (Boutilier, 1995). For each action A and state s , let $EFF(A, s)$ denote the

set of states s' such that $(s, A, s') \in R$. Hence $EFF(A, s)$ is the set of states that may result, given that action A is executed in state s .

Definition 12 *An effect ranking function is a function δ that maps every action-state pair (A, s) to a plausibility function over $EFF(A, s)$.*

Informally, an effect ranking function gives the plausibility of each possible effect for each action. For instance, if we want to model a coin flipping action, then the corresponding effect ranking function would be constant: both outcomes of the coin flip would be considered equally plausible.

A *non-deterministic graded world view* is a pair $\langle W, \delta \rangle$ where W is a graded world view and δ is an effect ranking function. We illustrate with an example.

Example Consider an action domain involving a single fluent symbol *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. We illustrate how non-deterministic graded world views can provide a representation of this problem.

Both actions have non-deterministic effects in that both may cause *LampOn* to become true, but both may also fail to do so. We define a graded world view $\langle ACT, OBS \rangle$ of length 1. First, we define *ACT* so that *ThrowPaper* is the most likely action at time 1.

	λ	<i>Press</i>	<i>ThrowPaper</i>
ACT_1	10	2	1

Next we define *OBS* so that initially the light is off, and then the light is on.

	\emptyset	<i>LampOn</i>
OBS_0	10	0
OBS_1	0	10

Finally, we define an effect ranking function δ that represents the fact that pressing is more likely to turn the light on.

	\emptyset	$\{LampOn\}$
$\delta(Press, \{LampOn\})$	0	10
$\delta(Press, \emptyset)$	10	0
$\delta(ThrowPaper, \{LampOn\})$	9	0
$\delta(ThrowPaper, \emptyset)$	0	9

Lines 1 and 2 in the table indicate that pressing the switch is very likely to change the state of the lamp. Lines 3 and 4 indicate that throwing paper at the switch is most likely

to have no effect, but a change caused by the paper ball is seen as more plausible than a failed effect when the button is pressed.

In the preceding example, there are two possible solutions: either a plausible event occurs with an unlikely outcome, or a less plausible event occurs with an expected outcome. There is no a priori preference given to occurrence plausibilities or to effect plausibilities; the framework is flexible enough to represent either possibility.

Introducing effect ranking functions makes the distinction between action occurrences and action effects explicit, which in turn gives a straightforward treatment of failed actions. However, we need to introduce some extra machinery in order to determine the most plausible action history. The most general approach is to extend the definition of an aggregate plausibility function: a *non-deterministic aggregate plausibility function* takes a non-deterministic graded world view as an argument, and it returns a plausibility function over histories. An *admissible* non-deterministic aggregate plausibility function is one that increases monotonically with respect to the given graded world view, as well as the given effect ranking function.

We have been using the function *sum* as our standard aggregate plausibility function. The natural extension of *sum* to non-deterministic graded world views is the following. For any history $h = s_0, A_1, \dots, A_n, s_n$, define

$$\text{sum}(\langle ACT, OBS \rangle, \delta)(h) = \sum_i OBS_i(s_i) + ACT_i(A_i) + \delta(A_i, s_{i-1})(s_i).$$

It is easy to see that this is an admissible non-deterministic aggregate function. Returning to the lamp example, there are two minimally ranked histories under this function: one in which the lamp was turned on by pressing on the switch and one in which the lamp was turned on by throwing a piece of paper at the switch.

In the remainder of this section, we will assume that *sum* is the default aggregate function for non-deterministic world views. Under this assumption, we demonstrate that non-deterministic graded world views can be translated into graded world views in an extended action signature.

Let $T = \langle S, R \rangle$ be a non-deterministic transition system over the action signature $\langle \mathbf{A}, \mathbf{F} \rangle$. Let $\langle \langle ACT, OBS \rangle, \delta \rangle$ be a non-deterministic graded world view. We extend the action signature to a new action signature \mathbf{A}' where every edge in T corresponds to a unique action symbol. In particular, let $\mathbf{A}' = \{A_{(s,A,t)} \mid (s, A, t) \in R\}$. Let $T' = \langle S, R' \rangle$ where R' is the closure of the set $\{\langle s, A_{(s,A,t)}, t \rangle \mid s, t \in S\}$. Suppose that $ACT = ACT_1, \dots, ACT_n$. Define $ACT' = ACT'_1, \dots, ACT'_n$ where, for each i , $ACT'_i(A_{(s,A,t)}) = ACT_i(A) + \delta(A, s)(t)$.

Proposition 5 *For any non-deterministic transition system T , a history*

$$h = s_0, A_1, \dots, A_n, s_n$$

obtains the same rank in $\langle \langle ACT, OBS \rangle, \delta \rangle$ as

$$h' = s_0, A_{(s_0, A_1, s_2)}, \dots, A_{(s_{n-1}, A_n, s_n)}, s_n$$

obtains in $\langle ACT', OBS \rangle$.

Proof The plausibility of h is obtained by taking the sum

$$\sum_i OBS_i(s_i) + ACT_i(A_i) + \delta(A_i, s_{i-1})(s_i),$$

which is clearly the same sum taken to determine the plausibility of h' . \square

Hence non-deterministic actions and failed actions can be represented in a graded world view, simply by setting up the plausibility functions carefully.

We remark that there is a conceptually interesting distinction that is lost in this translation. Informally, there is a distinction between an action that fails to occur and an action that occurs, but fails to produce an expected effect. This distinction is clear if we consider the difference between failing to drop a glass on the ground, and dropping a glass that fails to break when it hits the ground. In the first case, the agent executes the drop action but it fails to occur; perhaps the glass sticks to the agent's hand. In the second case, the glass is successfully dropped without breaking. In our framework, both of these events are represented by a dropping action with the null effect. We suggest that this is an acceptable treatment, because in both cases the sequence of actions and states is identical. As such, we cannot distinguish between these scenarios based on our definition of a history. However, we may be able to distinguish indirectly based on the values of other fluents. For instance, the location of the glass is only going to change in the case where it is successfully dropped.

5. Comparison with Related Formalisms

Belief change due to actions and observations has been addressed previously in the literature. In this section, we consider how our approach is related to existing work in the area.

5.1 Representing Single-Shot Belief Change

While our focus thus far has been on the use of graded world views for the representation of iterated belief change due to actions and observations, the simplest scenario involves just a single ontic or epistemic action. It is important, therefore, to verify that the single-shot belief change operators induced by a graded world view are reasonable with respect to existing work in the area. Recall that we defined $*$ and \diamond on graded world views as shorthand notation for the associated concatenation operations. Based on the results in this section, it will be clear that this shorthand is natural and appropriate.

We first consider the case of a single ontic action.

Proposition 6 *Let $\langle ACT, OBS \rangle$ be a graded world view. For any plausibility function r over \mathbf{A} ,*

$$Bel(\langle ACT, OBS \rangle \bullet \langle r, 0 \rangle) = Bel(\langle ACT, OBS \rangle) \diamond Bel(r).$$

Proof Since every action is always executable, $s \in Bel(\langle ACT, OBS \rangle \bullet \langle r, 0 \rangle)$ if and only if there is some $s' \in Bel(\langle ACT, OBS \rangle)$ such that $s' \diamond A = s$ for some $A \in Bel(r)$. Hence $s \in Bel(\langle ACT, OBS \rangle \bullet \langle r, 0 \rangle)$ if and only if $s \in Bel(\langle ACT, OBS \rangle) \diamond Bel(r)$. \square

Proposition 6 is important if we are primarily interested in belief states and ontic actions. Basically, in this case, graded world views are unnecessary. The most plausible final belief state can be determined by simply looking at the belief state associated with the initial graded world view.

We now consider the case of a single observation. At present, we are primarily interested in comparing the expressive power of graded world views with AGM revision operators. There is one sense in which graded world views are clearly more expressive than AGM operators. In particular, a new observation need not be incorporated into an agent's beliefs if the observation does not come from a reliable source. We will demonstrate that, in the context of a single observation, this is essentially the only difference between a graded world view and an AGM revision operator. More specifically, we will see that the belief change defined by concatenating a single observation onto a graded world view can be captured by an AGM operator, provided that the observation has degree of strength higher than some fixed threshold.

We start by proving that every plausibility function defines a system of spheres, as defined by Grove (1988). We first review the definition of a system of spheres. Let $M_{\mathcal{L}}$ denote a set of consistent, complete theories over \mathcal{L} . A set of subsets \mathcal{S} of $M_{\mathcal{L}}$ is a system of spheres centered on X where $X \subseteq M_{\mathcal{L}}$, if it satisfies the conditions:

- S1. \mathcal{S} is totally ordered by \subseteq
- S2. X is the minimum of \mathcal{S} under \subseteq
- S3. $M_{\mathcal{L}} \in \mathcal{S}$
- S4. For any formula ϕ with $|\phi| \neq \emptyset$, there is a least sphere $c(\phi)$ such that $c(\phi) \cap |\phi| \neq \emptyset$ and $U \cap |\phi| \neq \emptyset$ implies that $c(\phi) \subseteq U$ for every $U \in \mathcal{S}$

We picture a system of spheres as a series of concentric circles, with innermost circle X .

Let r be a plausibility function over X with minimum value \min_r . For any n , let $r[n]$ denote the set of complete, consistent theories that are satisfied by some interpretation I with $r(I) \leq n$.

Proposition 7 *Let r be a plausibility function over a finite action signature. The collection $\mathcal{R} = \{r[n] \mid n \geq \min_r\}$ is a system of spheres centered on $r[\min_r]$.*

Proof Clearly, for each n , $r(n) \subseteq r(n+1)$. Hence \mathcal{R} is totally ordered by \subseteq .

If $T \in r[\min_r]$, then T is satisfied by some I with $r(I) \leq \min_r$. But then, for any n , T is satisfied by some I with $r(I) \leq n$. Hence $r[\min_r] \subseteq r[n]$ for all $r[n]$.

Since the action signature is finite, there are only finitely many states. Hence there is a state that is assigned a maximum plausibility, say \max_r . Therefore, $r[\max_r]$ is the set of complete, consistent theories.

Let ϕ be a consistent formula. Since there are only finitely many states, there must be a state $s \in \kappa$ such that $r(s) \leq r(t)$ for all $t \in \kappa$. Let $n = r(s)$. Clearly $r(n) \cap \kappa \neq \emptyset$. Now suppose that $U \in \mathcal{S}$ and $U \cap \kappa \neq \emptyset$. Suppose that $U = r(m)$, so U is the set of complete, consistent theories satisfied by some I with $r(I) \leq m$. Since some elements of U are also in κ , it follows $m \geq n$. Therefore $r[n] \subseteq U$, and $r[n]$ is the least sphere intersecting κ . \square

Using this result, we can show that single-shot revision under graded world views can be captured by AGM revision operators. We make this claim precise in the next proposition.

Proposition 8 *Let $\langle ACT, OBS \rangle$ be a graded world view. There is an AGM revision function $*$ and a natural number n such that, for any plausibility function r over states with degree of strength larger than n ,*

$$Bel(\langle ACT, OBS \rangle \bullet \langle \lambda \uparrow n, r \rangle) = Bel(\langle ACT, OBS \rangle) * Bel(r).$$

Proof Recall that *plaus* is a plausibility function over histories that is defined by minimizing sums over $\langle ACT, OBS \rangle$, and *plaus_state* is the corresponding plausibility function over final states.

Let n be a natural number such that $n > \text{plaus}(h)$ for every history h . Let r be a plausibility function with rank n . It follows that $s \in Bel(\langle ACT, OBS \rangle \langle \lambda \uparrow n, r \rangle)$ if and only if the following conditions hold:

1. $s \in Bel(r)$
2. *plaus_state*(s) is minimal among all states satisfying 1.

By Proposition 7, *plaus_state* defines a system of spheres centered on $Bel(\text{plaus_state})$. It follows from Grove's representation result (Grove, 1988) that there is an AGM revision function $*$ such that, for any observation α , $s \in Bel(\text{plaus_state}) * \alpha$ if and only if the following conditions hold:

1. $s \in \alpha$
2. *plaus_state*(s) is minimal among all states satisfying 1.

Setting $\alpha = Bel(r)$ gives the desired result. \square

Proposition 8 illustrates that, for a single observation, the most plausible worlds can be determined without considering the history of actions and observations. We can determine the most plausible worlds following an observation by simply abstracting a belief state from a graded world view, then performing AGM revision. It is easy to show that the converse is also true: every AGM revision operator can be represented by a graded world view. More precisely, we have the following result.

Proposition 9 *Let $*$ be an AGM revision operator and let κ be a belief state. There is a graded world view $\langle ACT, OBS \rangle$ with $Bel(\langle ACT, OBS \rangle) = \kappa$ and a natural number n such that, for every non-empty observation α ,*

$$\kappa * \alpha = Bel(\langle ACT, OBS \rangle \bullet \langle \lambda \uparrow n, r \rangle)$$

where r is any plausibility function over states where the minimal ranked elements α have degree larger than n .

Proof By Grove's representation result, $*$ can be captured by a system of spheres \mathcal{S} . Define the graded world view $\langle ACT, OBS \rangle$ such that \mathcal{S} is the system of spheres given by Proposition 7. Set n such that $n > \text{plaus}(h)$ for every history h . The result is immediate. \square

Taken together, Propositions 8 and 9 illustrate that graded world views are equivalent to

AGM revision if we restrict attention to a single observation with a sufficiently high degree of reliability. Hence, for single-shot belief change, the full expressive power of graded world views is unnecessary. For both ontic actions and observations, we can define the same belief change operations if we start with just a belief state. There is also a correspondence here with Nayak’s work on iterated revision (Nayak, 1994); if an observation is sufficiently plausible, then every state in that observation ends up being strictly more plausible than every other state.

5.2 Representing Conditionalization

Spohn uses ranking functions to define a different form of single-shot belief change called *conditionalization* (Spohn, 1988). The idea is that new evidence is presented as a pair (α, m) , where α is a set of states and $m \geq 0$; the value of m is an indication of the strength of the observation α . Informally, the conditionalization of r is a new function where the minimally ranked α -worlds receive rank 0 and the non- α worlds are all “shifted up” by m . In this section, we illustrate how conditionalization can be defined in terms of graded world views.

First, we define conditionalization formally. Let r be a plausibility function with $\min_r = 0$ and let α be a subset of the domain of r . Let $\min(\alpha)$ denote the minimum value $r(s)$ for $s \in \alpha$. Spohn defines the the plausibility function $r(\cdot|\alpha)$ over α as follows:

$$r(s|\alpha) = r(s) - \min(\alpha).$$

We call $r(s|\alpha)$ the α -part of r . The conditionalization of r , written $r_{(\alpha,m)}$, is the following plausibility function.

$$r_{(\alpha,m)}(s) = \begin{cases} r(s|\alpha) & \text{if } s \in \alpha \\ m + r(s|\tilde{\alpha}) & \text{if } s \notin \alpha \end{cases}$$

So the conditionalization of r is the α -part of R together with the $\tilde{\alpha}$ -part shifted appropriately.

We show that conditionalization can easily be represented by taking minimal sums over plausibility functions.

Definition 13 *Let r be a plausibility function over $2^{\mathbf{F}}$, let α be a non-empty subset of $2^{\mathbf{F}}$, and let m be a natural number. Define $r_C(\alpha, m)$ as follows:*

$$r_C(\alpha, m)(s) = \begin{cases} 0 & \text{if } s \in \alpha \\ m + \min(\alpha) & \text{if } s \notin \alpha \end{cases}$$

We refer to $r_C(\alpha, m)$ as the *conditionalizer* of r with respect to α and m . The following proposition illustrates how we can define the conditionalization of a plausibility function by taking an appropriate sum.

Proposition 10 *Let r be a plausibility function with $\min_r = 0$. For any α, m , the normalization of $r + r_C(\alpha, m)$ is the conditionalization $r_{(\alpha,m)}$.*

Proof If $s \in \alpha$, then

$$r(s) + r_C(\alpha, m)(s) = r(s) + 0 = r(s).$$

If $s \notin \alpha$, then

$$r(s) + r_C(\alpha, m)(s) = r(s) + m + \min(\alpha).$$

Since $r(s) \geq 0$ and $m \geq 0$, it follows that the minimum value obtained by $r + r_C(\alpha, m)$ is $\min(\alpha)$. Hence, the normalization of $r + r_C(\alpha, m)$ is the plausibility function r' defined as follows.

$$r'(s) = r(s) + r_C(\alpha, m)(s) - \min(\alpha)$$

This is equal to $r_{(\alpha, m)}$, which is what we wanted to show. \square

Proposition 10 illustrates that the conditionalization of r by (α, m) can be defined by taking a minimal sum over two plausibility functions. We have restricted attention to plausibility functions with minimum 0 because this class coincides more closely with Spohn's ranking functions. However, we can define the conditionalizer in the same manner for plausibility functions with non-zero minimums. We can also define the conditionalization of a graded world view. Informally, we simply conditionalize the associated plausibility function on states. Hence, we identify the conditionalization with respect to $\langle \alpha, m \rangle$ with the following operation:

$$\langle ACT, OBS \rangle \bullet \langle \lambda, plaus_state_C(\alpha, m) \rangle.$$

It is straightforward to show that this gives the desired result.

5.3 Representing Belief Evolution Operators

As noted previously, we have previously defined so-called *belief evolution operators* to capture iterated belief change due to actions and observations (Hunter & Delgrande, 2011). In this section, we show that graded world views actually extend this approach, by verifying that every belief evolution operator can be captured by a graded world view.

There are two underlying assumptions in the definition of belief evolution:

1. The plausibility of an observation is determined by some ordering, *recency* by default.
2. The action history is assumed to be correct.

Both of these assumptions can be represented in a graded world view by setting up the plausibility functions appropriately.

Belief evolution operators can be defined with respect to metric transition systems. A *metric transition system* to be a transition system, along with a *metric* d that gives a distance between states. The metric d defines a belief revision operator as follows (Delgrande, 2004):

$$\begin{aligned} \kappa * \alpha &= \{w \in \alpha \mid \exists v_1 \in \kappa \text{ such that for all } v_2 \in \alpha, v_3 \in K \\ &\quad \text{we have } d(w, v_1) \leq d(v_2, v_3)\}. \end{aligned}$$

Assume that we have a fixed initial belief state κ_I , along with a metric transition system defining a revision operator $*$ and an update operator \diamond . Let \circ be the belief evolution operator obtained from $*$ and \diamond . Let

$$\bar{A} = \langle A_1, \dots, A_n \rangle$$

be an action trajectory, and let

$$\bar{\alpha} = \langle \alpha_1, \dots, \alpha_n \rangle$$

be an observation trajectory. We want to construct a graded world view W_{ev} that assigns minimal plausibility value to all histories corresponding to $\kappa_I \circ \langle \bar{A}, \bar{\alpha} \rangle$.

We define $W_{ev} = \langle ACT, OBS \rangle$ presently. By combining κ_I with the metric d given by the metric transition system, we can define a plausibility function $BASE$ that represents the initial ordering of states implicit in $*$. In particular, for any s , set

$$BASE(s) = \min(\{d(s, k) \mid k \in \kappa_I\}).$$

Using this plausibility function, we can define the observation trajectory OBS . Let \max denote the maximum value obtained by $BASE$.

$$OBS_i = \begin{cases} BASE & \text{if } i = 0 \\ \alpha_i \uparrow (2^i + \max) & \text{otherwise} \end{cases}$$

By incrementing the plausibility of false observations exponentially, we can assure that recent observations will be given greater credence.

Informally, each action symbol A_i is translated into a plausibility function that obtains the minimum value on the set $\{A_i\}$. Formally, we have the following, for $1 \leq i \leq n$:

$$ACT_i = A_i \uparrow (2^{n+1} + \max).$$

Proposition 11 *If $\kappa_I \circ \langle \bar{A}, \bar{\alpha} \rangle = \langle \kappa_0, \dots, \kappa_n \rangle$, then*

$$\begin{aligned} h &\in \Phi(W_{ev}) \\ &\iff \\ h &= \langle s_0, A_1, \dots, A_n, s_n \rangle \text{ where } s_i \in \kappa_i \text{ for each } i. \end{aligned}$$

Proof Assume for the moment that $\langle \bar{A}, \bar{\alpha} \rangle$ is consistent. Let $h = \langle v_0, B_1, \dots, B_n, v_n \rangle$. By definition $h \in \Phi(W_{ev})$ if and only if the sum

$$\sum_{i=1}^n ACT_i(B_i) + \sum_{i=0}^n OBS_i(v_i) \tag{2}$$

is minimal. Since $\langle \bar{A}, \bar{\alpha} \rangle$ is consistent, there exist histories $\langle s_0, A_1, \dots, A_n, s_n \rangle$ where each $s_i \in \alpha_i$. For such histories, the sum (2) becomes

$$\sum_{i=1}^n ACT_i(A_i) + \sum_{i=0}^n OBS_i(s_i) = 0 + OBS_0(s_0)$$

We remark that this sum is less than any sum that can be obtained by a history where there is some i such that either $B_i \in A_i$ or $s_i \notin \alpha_i$. Therefore $h \in \Phi(W_{ev})$ if and only if the following three conditions hold:

1. $B_i = A_i$ for each $i > 0$

2. $v_i \in \alpha_i$ for each $i > 0$
3. $OBS_0(v_0)$ is minimal among states satisfying 1 and 2.

In order to satisfy condition 2, it must be the case that v_0 is in the set

$$V = \bigcap_i \alpha_i^{-1}(\bar{A}_i).$$

In order to simultaneously satisfy condition 3, it must also be the case that v_0 is minimally distant from κ_I according to the metric d . In other words, $v_0 \in \kappa_I * V$. Therefore, $h \in \Phi(W_{ev})$ if and only if each $B_i = A_i$ and the following conditions hold:

1. $v_0 \in \kappa_I * \bigcap_i \alpha_i^{-1}(\bar{A}_i)$, and
2. $v_i = v_0 \diamond \bar{A}_i$.

This is the definition of $\kappa_I \circ \langle \bar{A}, \bar{\alpha} \rangle$, so this completes the proof.

The case where $\langle \bar{A}, \bar{\alpha} \rangle$ is inconsistent is similar. The only difference is that we need to notice that the degree of strength of each observation increases by a power of 2. We use the fact that, for any natural number p , 2^p is larger than every sum of terms 2^i with $i < p$. As such, in order to minimize the sum (2), we need to work backwards through the observations, keeping each observation if it is consistent with the observations that followed. This is just an equivalent specification of $\tau(W_{ev})$ – increasing powers exponentially forces a strict preference for recent observations. \square

Proposition 11 demonstrates that graded world views can represent any belief evolution operator defined with respect to a distance function. From the perspective of graded world views, the assumption that action histories are infallible is essentially just a restriction on the admissible plausibility functions.

We conclude this section with some brief remarks about the use of orderings to resolve inconsistency in iterated belief change. The well-known Darwiche-Pearl postulates for iterated belief change (Darwiche & Pearl, 1997) are only satisfied when we assume that the most recent observation takes precedence over previous observations. By contrast, Papini illustrates an alternative approach to iterated revision in which earlier observations take precedence over later observations (Papini, 2001). More generally, we defined belief evolution operators with respect to an arbitrary total ordering over the observations. The most natural extension of belief evolution would extend the ordering to include all observations and actions. Using the techniques in this section, it is easy to see that this extended conception of belief evolution corresponds to the class of graded world views with an arbitrary initial observation followed by plausibility functions of the form $\alpha \uparrow 2^i$, where each i is distinct. Hence, even the most general extension of belief evolution can be represented by a relatively restricted class of graded world views.

5.4 Relation with the Situation Calculus

The Situation Calculus (SitCalc) is a well-establishing framework for reasoning about the effects of actions. Action descriptions in the SitCalc are formulated in many-sorted first-order logic along with a single second-order induction axiom. Briefly, variables can range

over *situations*, *entities*, and *actions*. Situations are understood to represent the current state of the world, including a complete history of all actions that have occurred. There is a distinguished constant symbol S_0 that denotes the *initial situation* and a distinguished function symbol *do* that maps an action A and a situation s to the situation that results from executing A in situation s . Predicate symbols that take a situation argument are called *fluents*. For a detailed introduction, we refer the reader to the foundational summary presented by Levesque, Pirri and Reiter (1998).

The epistemic extension of the SitCalc introduces sensing actions, alternative initial situations, an accessibility relation, and a plausibility ordering over situations (Shapiro et al., 2011). The situations that an agent believes possible are the minimal accessible situations, and the effect of a sensing action changes the set of accessible states. This results in a form of belief revision that satisfies most of the AGM postulates. We have previously proved that belief revision in the SitCalc can be captured by a belief evolution operator under a natural translation (Hunter & Delgrande, 2011); it therefore follows from Proposition 11 that the belief change operators defined in the standard SitCalc approach can be captured by graded world views.

A similar conclusion can be drawn regarding our representation of uncertain actions. Bacchus et. al. (1999) extend the SitCalc to include “noisy” sensors and non-deterministic actions. Roughly, the idea is to define complex actions in terms of a set of *primitive actions*. In this manner, a non-deterministic action is represented by a probability function over the set of primitive actions. We have a similar result in Proposition 5, which proves that we can capture non-deterministic actions by introducing ranking functions over action effects. If we think of each action effect as a *primitive action*, then we essentially have the same representation as Bacchus et. al. except that we use ranking functions rather than probabilities. Moreover, in section 3.4 we demonstrated that a translation from probabilities to plausibilities is straightforward.

While our approach appears to be sufficiently general to capture belief change and uncertainty over actions as defined in the given epistemic extensions of the SitCalc, it is important to note that there are respects in which our approach is actually more expressive. In particular, we allow revision by arbitrary rankings over formulas; we do not require a fixed set of sensing actions. Our approach also allows for a flexible representation of the interaction between uncertainty over actions and observations, and it is elaboration tolerant in the sense that we need only modify aggregates over plausibility functions to capture different phenomena. On the other hand, the SitCalc provides a more rigorous and precise treatment of actions effects and ontic change. For example, the SitCalc treatment of *property persistence* (Kelly & Pearce, 2010) captures a natural feature of certain world properties that can only be captured by a graded world view with an ad hoc restriction on plausibility functions. Overall, we see our approach as a high-level model of belief change that can capture the dynamics of belief in the SitCalc, but it does so while losing some of the advantages of the SitCalc model of action effects.

5.5 Relation with Dynamic Epistemic Logic

The notion of belief change due to reported information has been addressed in Dynamic Epistemic Logic (DEL) (van Ditmarsch, van der Hoek, & Kooi, 2007). In this context, belief

change is captured through *plausibility models*. A plausibility model is a Kripke structure, where we associate a well-ordering over possible states with each state; this ordering indicates which states are seen as the most plausible at each state in the structure. Belief update can be captured in this setting through state-changing actions that allow different agents to have different awareness of the action that has been executed (Baltag & Smets, 2008); belief revision can be captured by a mapping on plausibility models that is syntactically defined through suitable modal operators (Baltag & Smets, 2006; Van Benthem, 2007). Although it is common to use an ordering to represent plausibility in DEL, there are also variants in which plausibility is captured with quantitative values. For example, Laverny and Lang (2005) define a logic that incorporates actions and uses ranking functions over $\mathbf{N} \cup \{\infty\}$ to model strength of belief for formulas without nested modalities.

There are two main distinctions between our work and related work in the DEL tradition. The first distinction is superficial, and it is related to our high-level perspective on new information. In DEL, there is an *external* perspective on new information, in which structures explicitly model the way that each agent in the system views the new information. This is why DEL is so expressive with respect to nested multi-agent beliefs: we are able to be explicit about the perspective of each agent. As presented in this paper, we can think of graded world views as taking an *internal* perspective for a single agent. When we provide a ranking function that represents an observation, we have no external knowledge about the “actual” state of the world. The advantage of this approach is that every observation (resp. action) is always possible. By contrast, with the external perspective of DEL, there may be situations where certain information simply cannot be provided as an information update. For example, if the underlying Kripke structure only contains states where ϕ is true, then there is no information update that can be provided to make an agent believe ϕ is false.

The second distinction between our work and existing work in DEL is related to our uniform representation of observations and actions. In DEL, a distinction is often drawn between so-called *hard updates* and *soft updates* (Van Benthem, 2007). A hard update occurs when the information is definite in some sense, and it is incorporated by modifying the accessible states. This might occur, for example, when it is common knowledge that a certain state-changing action has occurred. A soft update corresponds to a belief revision, where the information is incorporated by modifying the plausibility ordering. However, in a practical setting, it is often the case that one can not distinguish which kind of change should actually occur. This kind of ambiguity does not occur in the case of graded world views, where all belief change operations are handled in a uniform manner by taking aggregates over plausibility functions. Moreover, in both hard update and soft update, the update is typically by a “single” piece of information with a single level of plausibility. Hence, observations and actions are not represented in a graded manner. This is one of the distinguishing features of the present work: we allow observations (and actions) to provide different levels of evidence for several different states simultaneously.

Despite the differences, it is important to emphasize that we actually do not see graded world views as an alternative to DEL. Instead, we see the approach presented here as a tool that could be used as the basis for belief change in a suitably defined DEL with quantitative plausibility rankings. This is analogous to the way that formal belief change operators have been incorporated in modal logics. Iterated revision operators such as lexicographic revision (Nayak, 1994) and natural revision (Boutilier, 1996) were originally defined in a general

setting of orderings, before they were used as motivating examples for information update in DEL. Similarly, we suggest that a modal logic could be defined in which the semantics of iterated propositional updates was based on aggregates over plausibility functions. Such a logic would be able to model all of the practical domains discussed in this paper, and it would also be able to capture complex multi-agent situations involving nested belief.

5.6 Representations of Plausibility

Following Spohn (1988), we have taken quantitative ranking functions as our representation of plausibility. However, the literature on belief change and defeasible reasoning contains a great deal of work on different representations of plausibility. In this section, we briefly consider alternative approaches from the literature, and we argue that our approach is more suitable for the kind of reasoning we intend to capture.

The most basic representation of plausibility in the literature is a total pre-order on states, as we often see in AGM-inspired work. We normally think of such an ordering in terms of a series of “levels” of plausibility. The obvious limitation of this approach is that there can be no notion of magnitude with respect to differences in plausibility, although this can be introduced by allowing empty levels. In any event, it is clear that ranking functions are more expressive than total pre-orders with respect to modelling plausibility. A more interesting distinction occurs between our notion of plausibility, and the orderings used in the DEL tradition. As noted above, plausibility in this context is typically captured by a *well-ordering* (Van Benthem, 2007). This is necessary because, in this setting, we do not want all states to be comparable; we want certain states to be inaccessible from others. One might ask if our ranking functions are limited in any sense with respect to the well-orderings used in plausibility models. Our choice of ranking functions is motivated by the internal-external distinction discussed in the previous section. An individual agent need not consider “impossible” states in order to modify their beliefs appropriately, they need only consider highly unlikely states.

Note that our approach is superficially similar to related work on defeasible reasoning featuring some form of preferential ordering on states. It is possible, for example, to reason about typicality by supplementing a Kripke structure with a single ordering over states (Britz & Varzinczak, 2013). An alternative model in the setting of description logics has also been proposed, in which a function is introduced to map each concept to a set of elements that are deemed to be most typical (Giordano, Olivetti, Gliozzi, & Pozzato, 2013). However, despite the superficial similarities, it is important to be clear that our initial plausibility function is not intended to be a representation of “typicality” or “normality” in an absolute sense. Our model of plausibility corresponds more closely to a pointed plausibility model that includes a distinguished *actual state*.

Alternative quantitative models of plausibility have also been explored in the literature. One of the most influential such approaches is based on *possibility measures*, which are functions that map states to $[0, 1]$ (Benferhat et al., 1999). The main difference between a possibility measure and a plausibility function is really that the possibility measure 0 actually corresponds to impossibility, and there is no corresponding plausibility value. However, we have already shown that plausibility functions can be used to capture rational-valued

functions over $[0,1]$; it would be straightforward to modify our proof to deal with real number values.

One of the most general measures of plausibility in the literature is a *plausibility space*, which is analogous to a probability space except that the range of plausibility values can be any set (Friedman & Halpern, 2001). Plausibility spaces are intended to be a very general model that subsumes a great deal of existing work on reasoning about plausibility. Indeed, plausibility functions are clearly just a special case of a plausibility space; so our representation of belief can be seen as a special case of this more general approach.

In this section, we have been restricting the discussion to alternative representations of plausibility on states. As noted previously, our approach to the representation of plausibility over states is actually not new and it is therefore easy to position our work in this context. The novel feature of our approach is that we use the same notion of plausibility over actions, in order to develop a novel approach to reasoning about iterated belief change due to actions and observations

6. Limitations and Advantages

Our focus in previous sections has been on establishing the expressive power of graded world views, as compared with existing frameworks for reasoning about belief change. As a result, we have focused on problems involving an a priori graded world view, along with some “new” information. However, the restriction to new information is artificial. In the general case, there is no reason to restrict attention to problems in which an agent only receives information about actions and observations occurring at the most recent point in time. An agent could certainly receive new information about earlier events and actions. Hence, a more general problem involves an agent with an underlying graded world view, together with a set of constraints on the most plausible histories. In this section, we consider the representation of problems that have this more general form.

6.1 Constrained World Views

Suppose that $\langle ACT, OBS \rangle$ is a graded world view of length n . An *action constraint* is a pair (A, i) where A is an action symbol and $i \leq n$. Define $\Phi(\langle ACT, OBS \rangle) \upharpoonright (A, i)$ to be the set of histories with minimal plausibility, subject to the restriction that the i^{th} action executed is A . We define *observation constraints* in the analogous manner, and we let $\Phi(\langle ACT, OBS \rangle) \upharpoonright (\alpha, i)$ be the set of minimally ranked histories where the i^{th} state is in α . If Ω is a set of constraints, then we define $\Phi(\langle ACT, OBS \rangle) \upharpoonright \Omega$ to be the set of minimally ranked histories satisfying every constraint in Ω . We will refer to such histories as *constrained histories* and we will refer to a graded world view together with a set of constraints as a *constrained world view*.

We have presented constrained world views to illustrate that graded world views are useful for many problems beyond those that are normally considered to be in the realm of a standard “belief change operator.” For example, suppose that Bob sends an encrypted email message to Alice, inviting her to a party at his house. Bob is aware that Eve is the system administrator, and that she could potentially manipulate the message before delivering it. When Alice does not show up, Bob concludes that Eve did not deliver the message. Bob is concerned that Eve read the message and had hurt feelings that she was

not invited. However, looking at every possible action Eve could take, Bob concludes that Eve could not have decrypted the message.

In the preceding example, Bob needs to consider all possible actions that Eve could have executed. The conclusion that Bob draws is that Eve’s knowledge of the party is invariant with respect to her actions. We can formally define invariance as follows.

Definition 14 *Let $\langle ACT, OBS \rangle$ be a graded world view. We say that a set of states α is an i -invariant of $\langle ACT, OBS \rangle$ if and only if, for every $A \in \mathbf{A}$, $Bel(\langle ACT, OBS \rangle \upharpoonright (A, i)) \subseteq \alpha$.*

The intuition behind i -invariance is that, regardless of the action at time i , the underlying agent will always believe that the actual world is in α . Reasoning about invariant properties is essential if an agent is trying to ensure some property must hold in an action domain involving exogenous actions. This is required, for example, in reasoning about cryptographic protocols.

Reasoning about invariance is just one new kind of problem that can be addressed by constrained world views. We suggest that constraints can also be used to provide natural representations of hypothetical reasoning and abductive reasoning.

6.2 Belief Extrapolation

Constrained world views are similar to *belief extrapolation operators* (Dupin de Saint-Cyr & Lang, 2011). Briefly, a belief extrapolation operator \uparrow takes a sequence of formulas, called a *scenario* as input, and it outputs another scenario. The intuition is that the output gives the most general sequence of formulas that can possibly be true, given the input and the assumption that fluents tend to be inertial. We give a brief description of the basic construction.

A *trajectory* is a sequence τ of interpretations over some fixed signature. Let $\tau(i)$ denote the i^{th} interpretation in the trajectory τ . Given a scenario Σ , let $Traj(\Sigma)$ denote the set of trajectories that satisfy each formula in Σ on a point-by-point basis. Every ordering \preceq on the class of trajectories defines an extrapolation operator \uparrow as follows:

$$|(\Sigma \uparrow (t))| = \{\tau(t) \mid \tau \in Min(\preceq, Traj(\Sigma))\}$$

Hence, $\Sigma \uparrow$ picks out the minimal trajectories satisfying Σ .

For our present purposes, the most important feature of a belief extrapolation operator is that it is defined with respect to an ordering over histories. Given an ordering over histories together with a sequence of formulas, a belief extrapolation operator returns the most plausible sequences of states. In the case of constrained world views, we essentially do the same thing. The given graded world view defines an ordering over states, and the constraints give a sequence of conditions that need to be satisfied. It is easy to show that, using the summation aggregate, every graded world view can be captured by a belief extrapolation operator. In this section, we consider the converse problem: Can every belief extrapolation operator be represented by a suitable graded world view?

The key observation here is that the mapping from graded world views to orderings on histories is not surjective; there are orderings on histories that can not be described by a graded world view. For example, a graded world view can not explicitly capture plausibilities of the form “if A_1 occurs at time i , then A_2 is likely to occur at time $i + 1$.”

This is easy to see, because a graded world view ranks histories by combining a sequence of rankings at each instant. Of course, it might be possible to devise a sequence of rankings along with an aggregate function that happened to support a conditional plausibility for a fixed action signature. But this would be a domain-specific property, because the temporal relation between actions can not be expressed with two independent plausibility functions. Informally, a graded world view can only represent domains where the ordering on histories is built up in a pointwise manner by the plausibilities at each point in time. In this section, we use this limitation to establish a difference in expressive power between constrained world views and belief extrapolation operators.

First, we need to formalize the problem that we would like to address more precisely. Given a belief extrapolation operator, we would like to be able to find a graded world view that captures the same information.

Definition 15 *Let \downarrow be a belief extrapolation operator. We say that \downarrow is representable if there is a graded world view $\langle ACT, OBS \rangle$ such that, for every scenario Σ of length n ,*

$$Traj(\Sigma \downarrow) = \Phi(\langle ACT, OBS \rangle) \upharpoonright \Sigma.$$

If \downarrow is representable, then the behaviour of \downarrow can be simulated with a graded world view. We remark that we have abused notation in the definition in that $Traj(\Sigma \downarrow)$ is a collection of sequences of states, whereas $\Phi(\langle ACT, OBS \rangle) \upharpoonright \Sigma$ is a collection of histories. We interpret the equality to mean that the two collections are equal if we ignore the action symbols in the latter.

The following proposition indicates that belief extrapolation operators have an expressive advantage.

Proposition 12 *There is a belief extrapolation operator \downarrow that is not representable.*

Proof Let \preceq be an ordering in which the following trajectories are minimal.

1. $\langle \{a, b\}, \{a, b\}, \{\neg a, b\} \rangle$
2. $\langle \{a, b\}, \{a, b\}, \{a, b\} \rangle$
3. $\langle \{a, \neg b\}, \{a, b\}, \{\neg a, b\} \rangle$

Let \downarrow be the associated belief extrapolation operator. We will show that \downarrow is not representable.

Let $\Sigma = \langle a, a \wedge b, b \rangle$. Note that Σ is satisfied by all three minimal trajectories. Therefore $Traj(\Sigma \downarrow)$ is precisely the set of minimal trajectories.

Now suppose that $\langle ACT, OBS \rangle$ is a graded world view such that

$$\langle ACT, OBS \rangle \upharpoonright \Sigma$$

assigns minimal plausibility to 1, 2, and 3. Hence, there exist actions $A_1, A_2, A_3, B_1, B_2, B_3$ such that the following sums all obtain the minimum possible rank:

1. $OBS_0(\{a, b\}) + ACT_1(A_1) + OBS_1(\{a, b\}) + ACT_2(B_1) + OBS_2(\{\neg a, b\})$
2. $OBS_0(\{a, b\}) + ACT_1(A_2) + OBS_1(\{a, b\}) + ACT_2(B_2) + OBS_2(\{a, b\})$

$$3. \text{OBS}_0(\{a, -b\}) + \text{ACT}_1(A_3) + \text{OBS}_1(\{a, b\}) + \text{ACT}_2(B_3) + \text{OBS}_2(\{-a, b\})$$

It must be the case that $\text{ACT}_1(A_1) = \text{ACT}_1(A_2)$, because otherwise either 1 or 2 could be reduced by changing the first action. Similarly, it must be the case that $\text{ACT}_2(B_1) = \text{ACT}_2(B_3)$, because otherwise either 1 or 3 would not be minimal. So, we can rewrite the sums as follows:

1. $\text{OBS}_0(\{a, b\}) + \text{ACT}_1(A_1) + \text{OBS}_1(\{a, b\}) + \text{ACT}_2(B_1) + \text{OBS}_2(\{-a, b\})$
2. $\text{OBS}_0(\{a, b\}) + \text{ACT}_1(A_1) + \text{OBS}_1(\{a, b\}) + \text{ACT}_2(B_2) + \text{OBS}_2(\{a, b\})$
3. $\text{OBS}_0(\{a, -b\}) + \text{ACT}_1(A_3) + \text{OBS}_1(\{a, b\}) + \text{ACT}_2(B_1) + \text{OBS}_2(\{-a, b\})$

From 1 and 2, it follows from basic algebra that

$$\text{ACT}_2(B_1) + \text{OBS}_2(\{-a, b\}) = \text{ACT}_2(B_2) + \text{OBS}_2(\{a, b\}).$$

Substituting this in 3 gives another minimal sum:

$$\text{OBS}_0(\{a, -b\}) + \text{ACT}_1(A_3) + \text{OBS}_1(\{a, b\}) + \text{ACT}_2(B_2) + \text{OBS}_2(\{a, b\}).$$

This corresponds to the trajectory

$$\langle \{a, -b\}, \{a, b\}, \{a, b\} \rangle.$$

Hence, any graded world view assigning minimum plausibility to 1-3, must also assign minimum plausibility to this fourth trajectory. Informally, if 1-3 are preferred trajectories according to a graded world view, then we are forced to accept another preferred trajectory. But we already saw that $\text{Traj}(\Sigma \uparrow)$ consists only of 1-3. Therefore \uparrow is not representable. \square

Note that the proof of Proposition 12 is constructive and it demonstrates that there is a simple, concrete, extrapolation operator that is not representable.

Informally, Proposition 12 follows from the fact that some orderings on histories can not be defined by a graded world view. This is particularly important in applications where an agent has preferences over the order in which events occur. In such applications, it can be useful to assign plausibilities to certain sequences of actions. We suggest, however, that the class of orderings definable by graded world views is a natural class of orderings. In particular, there are many action domains where an agent has no preconceived assumptions about the order that exogenous actions will occur. Graded world views provide a reasonable tool for the representation of such action domains. Graded world views also have an expressive advantage in that they can explicitly represent fallible actions as well as other forms of uncertainty over action effects.

We have seen that we can not represent every belief extrapolation operator, but it would be a mistake to conclude that graded world views are strictly less expressive. Belief extrapolation really says nothing explicit about actions, which leads to several differences from our approach. Given an ordering over histories of length n , it is not clear how to use belief extrapolation to incorporate a new action followed by a new observation; there is no fixed method for extending orderings over n -tuples to orderings over $n + 1$ tuples. In

the case of a graded world view, however, it is clear how the new ordering is defined when more actions are performed. As such, graded world views are more appropriate for the representation of epistemic action domains where we expect new observations and actions to occur. Of course, in our approach, we do have a fixed length for the initial graded world view. Since we can have null actions, however, the value n is really a maximum length for the sequence of actions that occurred. As such, graded world views are most sensible in domains where we can imagine the maximum number of actions that have previously occurred is bounded. Again, we suggest that the class of domains that satisfies this constraint is both large and natural.

6.3 Beliefs about Action Effects

Despite our flexible representation of observations and actions, there is one aspect of our representation that is quite rigid. Specifically, we have a single underlying transition system that gives the effects of actions; there is no mechanism for changing this transition system in response to new information. This has been addressed by Varzinczak (2010) in a modal setting, by introducing a formal mechanism that allows the effects of actions to change in response to new information. Combining this idea with the notion of plausibility in a graded world view would give an even more complete picture of the way that agents reason about the effects of actions with uncertain effects. We leave a treatment of this idea for future work.

A related problem is the fact that the underlying transition system in our framework is essentially known to the agent. In other words, given perfect information about the current state of the world and the action that has been executed, then the agent knows the resulting state. An alternative approach due to Eiter et al. (2010) is to start with an *action description* that specifies some actions effects. We can then introduce suitable information update mechanisms based only on this (partial) description of effects, rather than by assuming a definite underlying transition system. This is, however, only a superficial distinction. By using effect ranking functions, we can essentially give a partial description of action effects that corresponds to an action description. In this manner, it is straightforward to define so-called *Action Description Update* problems with graded world views.

7. Conclusion

We have introduced a formalism for reasoning about sequences of actions and observations. The formalism uses Spohn-style ranking functions at each instant to determine the most plausible action or observation, and determines the most plausible histories by an aggregate function over all instants. We have proved that the formalism subsumes belief revision, belief evolution, and conditionalization. Moreover, it is suitable for the representation of fallible beliefs, erroneous perception, exogenous actions, and failed actions. We have used transition systems for the representation of actions in order to facilitate comparison with a wide range of action formalisms. In future work, we will be interested in axiomatizing the belief change that is permitted by the class of admissible aggregate functions.

The main advantage of graded world views over related formalisms is that graded world views provide a uniform mechanism for dealing with imperfect information about actions and observations. As such, graded world views allow observations that need not be incor-

porated because the agent is convinced a certain action occurred previously. On the other hand, graded world views also allow an agent to retract the belief that an action occurred if they have a strong belief in some observation. Graded world views provide a tool for representing these opposing conclusions, simply by considering the magnitudes of plausibilities, as well as some underlying aggregate.

In existing work, belief change caused by actions is often represented by starting with an action formalism and then adding revision operators. One problem with this approach is that it does not allow beliefs about action occurrences. In a sense, graded world views take the opposite approach. We start with ranking functions, which were originally defined for reasoning about belief change in a static environment, and then we plug in actions. An agent's beliefs about the actions that occur are independent of the formal representation of action effects. As such, although we have presented graded world views in terms of transition systems, it would certainly be possible to use a different action formalism. The key point is that, by using ranking functions to represent uncertainty about states and actions, we can define a framework for reasoning about epistemic action effects in which primary importance is placed on the evolution of an agent's beliefs.

We conclude with a brief remark about the overall approach taken in our framework. A distinction is commonly drawn between update and revision in the belief change literature, despite the fact that it can be difficult in practice to determine which is the appropriate operation given a particular piece of information. In many cases, it simply is not clear if conflicting information is the result of some unseen action or if it is the result of erroneous beliefs. One solution to this problem is to define a single, general belief change operation that subsumes both (Kern-Isberner, 2008). By contrast, we maintain an explicit distinction between the way beliefs change due to actions and observations, but we use a uniform model of uncertainty at the agent level that focuses on finding the most plausible sequence of events to explain all available evidence. This seems like a natural approach in many applications, where an agent's perception of past events is likely to be influenced by their conviction with respect to their senses and beliefs.

References

- Alchourrón, C., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50(2), 510–530.
- Bacchus, F., Halpern, J., & Levesque, H. (1999). Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111(1-2), 171–208.
- Baltag, A., & Smets, S. (2006). Dynamic belief revision over multi-agent plausibility models. In *Proceedings of Logic and the Foundations of Game and Decision Theory (LOFT)*, pp. 11–24.
- Baltag, A., & Smets, S. (2008). A qualitative theory of dynamic interactive belief revision. In *Proceedings of Logic and the Foundations of Game and Decision Theory (LOFT)*.
- Benferhat, S., Dubois, D., & Prade, H. (1999). Possibilistic and standard probabilistic semantics of conditional knowledge bases. *Journal of Logic and Computation*, 9(6), 873–895.

- Boutilier, C. (1995). Generalized update: Belief change in dynamic settings. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI 1995)*, pp. 1550–1556.
- Boutilier, C. (1996). Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic*, 25, 263–305.
- Britz, K., & Varzinczak, I. (2013). Defeasible modalities. In *Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pp. 49–60.
- Darwiche, A., & Pearl, J. (1997). On the logic of iterated belief revision. *Artificial Intelligence*, 89(1-2), 1–29.
- Delgrande, J. (2004). Preliminary considerations on the modelling of belief change operators by metric spaces. In *Proceedings of the 10th International Workshop on Non-Monotonic Reasoning (NMR 2004)*, pp. 118–125.
- Delgrande, J., Dubois, D., & 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., & Levesque, H. (2012). Belief revision with sensing and fallible actions. In *The Thirteenth International Conference on Principles of Knowledge Representation and Reasoning (KR2012)*, pp. 148–157.
- Dupin de Saint-Cyr, F., & Lang, J. (2011). Belief extrapolation (or how to reason about observations and unpredicted change). *Artificial Intelligence*, 2, 760–790.
- Eiter, T., Erdem, E., Fink, M., & Senko, J. (2010). Updating action domain descriptions. *Artificial Intelligence*, 174(15), 1172–1221.
- Friedman, N., & Halpern, J. (2001). Plausibility measures and default reasoning. *Journal of the ACM*, 48(4), 648–685.
- Gelfond, M., & Lifschitz, V. (1998). Action languages. *Linköping Electronic Articles in Computer and Information Science*, 3(16), 1–16.
- Giordano, L., Olivetti, N., Gliozzi, V., & Pozzato, G. (2013). A non-monotonic description logic for reasoning about typicality. *Artificial Intelligence*, 195, 165–202.
- Grove, A. (1988). Two modellings for theory change. *Journal of Philosophical Logic*, 17, 157–170.
- Hunter, A. (2014). Belief change and non-deterministic actions. In *Proceedings of the Canadian Conference on Artificial Intelligence*, pp. 289–294.
- Hunter, A., & Booth, R. (2015). Trust-sensitive belief revision. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI15)*.
- Hunter, A., & Delgrande, J. (2006). Belief change in the context of fallible actions and observations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI06)*, pp. 257–262.
- Hunter, A., & Delgrande, J. (2011). Iterated belief change due to actions and observations. *Journal of Artificial Intelligence Research (JAIR)*, 40, 269–304.

- Jin, Y., & Thielscher, M. (2004). Representing beliefs in the fluent calculus. In *Proceedings of the European Conference on Artificial Intelligence (ECAI04)*.
- Katsuno, H., & 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)*, pp. 387–394.
- Katsuno, H., & Mendelzon, A. (1992). On the difference between updating a knowledge base and revising it. In Gärdenfors, P. (Ed.), *Belief Revision*, pp. 183–203. Cambridge University Press.
- Kelly, R., & Pearce, A. (2010). Property persistence in the situation calculus. *Artificial Intelligence*, 174(12-13), 865–888.
- Kern-Isberner, G. (2008). Linking iterated belief change operations to nonmonotonic reasoning. In *Proceedings of the Eleventh International Conference on Principles of Knowledge Representation and Reasoning (KR08)*, pp. 166–176.
- Kraus, S., Lehmann, D., & Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 4, 167–207.
- Laverny, N., & Lang, J. (2005). From knowledge-based programs to graded belief-based programs, part i: On-line reasoning. *Synthese*, 147(2), 277–321.
- Lehmann, D., & Magidor, M. (1992). What does a conditional knowledge base entail?. *Artificial Intelligence*, 55, 1–60.
- Levesque, H., Pirri, F., & Reiter, R. (1998). Foundations for the situation calculus. *Linköping Electronic Articles in Computer and Information Science*, 3(18), 1–18.
- Liberatore, P., & Schaerf, M. (2000). BReLS: A system for the integration of knowledge bases. In *Proceedings of KR2000*, pp. 145–152. Morgan Kaufmann Publishers.
- Lorini, E., Jiang, G., & Perrussel, L. (2014). Trust-based belief change. In *ECAI 2014 - 21st European Conference on Artificial Intelligence*, pp. 549–554.
- Moore, R. (1985). A formal theory of knowledge and action. In Hobbs, J., & Moore, R. (Eds.), *Formal Theories of the Commonsense World*, pp. 319–358. Ablex Publishing.
- Nayak, A. (1994). Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41, 353–390.
- Papini, O. (2001). Iterated revision operations stemming from the history of an agent’s observations. In Rott, H., & Williams, M. (Eds.), *Frontiers in Belief Revision*, pp. 279–301. Kluwer Academic Publishers.
- Shapiro, S., Pagnucco, M., Lesperance, Y., & Levesque, H. (2000). Iterated belief change in the situation calculus. In *Proceedings of the Seventh International Conference on Principles of Knowledge Representation and Reasoning (KR 2000)*, pp. 527–538. Morgan Kaufmann Publishers.
- Shapiro, S., Pagnucco, M., Lesperance, Y., & 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., & Skyrms, B. (Eds.), *Causation in Decision, Belief Change, and Statistics, vol. II*, pp. 105–134. Kluwer Academic Publishers.

- Van Benthem, J. (2007). Dynamic logic for belief revision. *Journal of applied non-classical logics*, 17(2), 129–155.
- van Ditmarsch, H., van der Hoek, W., & Kooi, B. (2007). *Dynamic Epistemic Logic*. Springer.
- Varzinczak, I. (2010). On action theory change. *Journal of Artificial Intelligence Research (JAIR)*, 37, 189–246.