Charles University in Prague, Faculty of Arts
Department of Logic

Bc. Alexandra Kuncová

Master's Thesis

Generic approach to updating uncertainty: Focus on conditioning

Generický prístup k zmene neistoty s dôrazom na kondicionalizáciu

Supervisor: PhDr. Michal Peliš, Ph.D.
Consultant: Guillaume Aucher, Ph.D.

2016

Acknowledgements

First and foremost, I would like to thank my consultant Guillaume Aucher, Ph.D. It was far from easy to communicate over such a long distance and with so little time at hand. Yet my thesis is literally built on his ideas and suggestions. Also I would like to thank my supervisor PhDr. Michal Peliš, Ph.D. for his guidance. Last but not least, I would like to thank my family and friends who stood by my side even at times when the road turned bumpy.

Prehlasujem, že som diplomovú prácu vypracovala samostatne, že som správne citovala všetky použité pramene a literatúru a že práca nebola využitá v rámci iného vysokoškolského štúdia či k získaniu iného alebo rovnakého titulu.

V Prahe 3. januára 2016

Bc. Alexandra Kuncová

Abstract

First, we consider different kinds of representation of uncertainty and the methods for updating each of them by conditioning. We focus on the generic framework of (conditional) plausibility spaces, since it generalises all the introduced representations. Further, we select three frameworks and list the properties that need to be added to a conditional plausibility space in order to recover each of these frameworks. The main goal of this work, however, is to show how public announcement on single-agent plausibility models, ranking structures, and possibility structures realised by their corresponding update mechanisms, can be embedded into the framework of conditional plausibility spaces. At the very end we briefly illustrate a general update model using plausibility measures.

Keywords: belief revision, dynamic logic, epistemic logic, plausibility space, public announcement, uncertainty, update.

Abstrakt

Najprv sa budeme zaoberať rôznymi druhmi reprezentácie neistoty a metódami aktualizácie pomocou kondicionalizácie. Zameriame sa na generický rámec (kondicionálnej) plauzibility pre jeho schopnosť generalizovať všetky ostatné reprezentácie. Následne sa lepšie pozrieme na tri štruktúry a uvedieme vlastnosti, ktoré je nutné pridať rámcu kondicionálnej plauzibility, aby sme každú z nich dokázali skonštruovať. Hlavným cieľom tejto práce je ukázať, ako môže byť verejné vyhlásenie na jedno-agentových modeloch plauzibility, „ranking“ štruktúrach, a „possibility“ štruktúrach, v spojení s ich odpovedajúcimi mechanizmami zmeny, vnorené do rámca kondicionálnej plauzibility. Na záver ešte stručne popíšeme všeobecný model zmeny postavený na miere plauzibility.

Kľúčové slová: dynamická logika, epistemická logika, neistota, rámec plauzibility, revízia presvedčenia, verejné vyhlásenie, zmena neistoty.

Contents

Introduction

Epistemic logic, a modal logic suitable for reasoning about knowledge, has received quite the attention since the 1950s–1960s (thanks to G.H. von Wright, and J. Hintikka, respectively). With a group of agents we are able to investigate their individual knowledge, their knowledge about other agents' knowledge, and the alternatives an agent considers possible or impossible. What is more, we can even study distributive and common knowledge within a group of agents, i.e., what would they know if they put all their individual pieces of information together, and what everybody knows that everybody knows that everybody knows. . . , respectively. With more expressive dynamic epistemic logic we can capture the social aspects between the agents, that is, how their knowledge (individual and about the others) changes when they interact with each other via some kind of action (the most common being communication and announcements, but also including other 'nonverbal' actions as, for example, seeing somebody doing something).

An agent might not always be able to decide between the alternatives, she simply does not know which ones are the case. However, she still might prefer some of the alternatives more than the others, and thus she might be able to make at least an assumption according to her beliefs. After all, that is what we, people, usually do in real life – we decide according to our beliefs much more often than according to a 'bullet-proof' piece of knowledge.

Actions do not necessarily lead to gaining knowledge, to learning. However, they might still lead an agent to make a weaker yet more common kind of change, namely, the change of beliefs. The agent can add new beliefs, remove some of the old ones, or even revise them, so that in contrast with the first case, all her beliefs stay consistent. If the agent inclines to one alternative and an event takes place suggesting that the other alternative might be the case, she might be willing to change her mind, revise her beliefs, and eventually decide for the other alternative.

The first two chapters are based on the overview in [Hal03]. In the first chapter we introduce different kinds of representations of uncertainty, some quantitative, some qualitative – including probability measures, ranking functions, possibility measures, Dempster-Shafer belief functions, and the one generalising all of them, namely, plausibility measures.

In the second chapter we show how updating of uncertainty works for each of the representations mentioned in the first chapter. As the last one we present a conditional plausibility space, which is (in a certain sense) a generalisation of all the others described in this chapter.

In the third and main chapter we build our results on this generic framework. Our goal is to show that being equipped with a conditional plausibility space with appropriate additional properties, using conditioning we can 'simulate' certain actions (e.g., public announcement and radical revision) on three other frameworks for Baltag and Smets' single-agent plausibility models (in [BaS08]), ranking structures, and possibility structures, respectively.

In the fourth chapter we present a general model for revision using plausibility measures as suggested in [MaL08]. Even though this topic is beyond the scope of this thesis, we consider it very relevant and useful for future research in this area.

Chapter 1

Representing Uncertainty: Different Kinds of Representation

1.1 Language, Structures and Semantics

Let us first define several basic notions which are used on regular basis throughout this thesis starting with the syntax.

Definition 1.1 (Doxastic language)**.** For any countable set of atomic propositions Φ, we define *doxastic language* $\mathcal{L}_\Phi$ by:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid B^\varphi \varphi,$$

where $p \in \Phi$ and B^φ is the conditional doxastic modality (conditional belief). We use the usual abbreviations for the other boolean connectives as well as for $\top$ and $\bot$, and the abbreviation B for $B^\top$.

The atomic propositions (usually labeled by letters such as p and q) can be intuitively thought of as representing statements about the basic facts of the situation which must be either true or false (e.g., "It is raining." "5 is a prime.").

The core idea for all the upcoming formalisms comes from the well-known Kripke semantics for modal logic.

Definition 1.2. Given a countable set of atomic propositions Φ and a finite set of agents G, a *Kripke model* is a structure $M = (W, R, \pi)$, where

1. W is a set of states. The set W is also called the domain $\mathcal{D}(M)$ of M.
2. R is a function, yielding for every $i \in G$ an accessibility relation $R(i) \subseteq W \times W$.[1]
3. $\pi : \Phi \to 2^W$ is a valuation function that for every $p \in \Phi$ yields the set $\pi(p) \subseteq W$ of states in which p is true.

[1] Indeed, here $\times$ denotes a Cartesian product: $A \times B = \{(a, b) : a \in A \text{ and } b \in B\}$.

Since the main focus of this thesis is on belief and not knowledge, we restrict ourselves to the simplified single-agent cases. Thus, we omit all the accessibility relations except for the one (which obviously must be universal $W \times W$, i.e., defined on all the pairs of worlds) and add a plausibility measure or a plausibility relation, that is, a certain way of expressing plausibility.

The states in W are usually interpreted as possible worlds. Intuitively, they are the worlds that an agent considers possible. For our purposes, the 'objects' that are known (or considered likely or possible or probable) are propositions. Formally, a proposition is just a set of possible worlds, where it is true, i.e., a proposition "it is snowing in Prague," would correspond to the set of possible worlds where it is snowing in Prague.

The set of possible worlds that an agent considers possible can be viewed as a qualitative measure of her uncertainty. The more worlds she considers possible, the more uncertain she is as to the true state of affairs, and the less she knows. Suppose that an agent's uncertainty is represented by a set of possible worlds W. She considers Upossible if $U \cap W \neq \emptyset$; and she knows U if $W \subseteq U$.

Let us define what it means for a propositional formula φ to be true at a world w in model M. As far as the modal operator B^{ψ} for conditional belief is concerned, we shall define the exact condition for each of the selected frameworks individually later on. However, the main idea is that "...[c]onditional beliefs 'pre-encode' beliefs that [an agent] would have if [she] learnt certain things" [Ben04, p. 11]. It means that if $B^{\psi}\varphi$ is true, then it is more likely that $\varphi \wedge \psi$ is the case rather than $\neg\varphi \wedge \psi$ (i.e., if ψ, then typically φ).

Definition 1.3 (Satisfaction Relation)**.** Let $M = (W, R, \pi)$ be a Kripke model. The *satisfaction relation* is given by, for $w \in W$, $p \in \Phi$, φ, $\varphi' \in \mathcal{L}_{\Phi}$,

$$
\begin{array}{lll}
M, w \models p & \text{iff} & w \in \pi(p) \\
M, w \models \neg\varphi & \text{iff} & M, w \not\models \varphi \\
M, w \models \varphi \wedge \varphi' & \text{iff} & M, w \models \varphi \text{ and } M, w \models \varphi'
\end{array}
$$

We write $M \models \varphi$ to mean $M, w \models \varphi$ for all $w \in W$. Further, $\models \varphi$ (φ is valid) means that $M \models \varphi$ for all models M.

Last but not least, we will base some of the frameworks on sets of worlds rather than on worlds themselves. For that purpose we define an algebra.

Definition 1.4. An *algebra over* W is a set $\mathcal{F}$ of subsets of W that contains W and is closed under union and complementation, so that if U and V are in $\mathcal{F}$, then so are $U \cup V$ and $\overline{U}$.

Note that an algebra is also closed under intersection, since $U \cap V = \overline{\overline{U} \cup \overline{V}}$.

Now we can proceed to the overview of frameworks based on [Hal03, ch. 2]. Naturally, we place more emphasis on those which will be used further on. The rest is mentioned to stress the generality of the plausibility measures.

1.2 Probability Measures

Perhaps the best-known approach to getting a more fine-grained representation of uncertainty is probability.

Definition 1.5. A *probability space* is a tuple $(W, \mathcal{F}, \mu)$, where W is a non-empty set of possible worlds, $\mathcal{F}$ is an algebra of measurable subsets of W (i.e., a set of subsets closed under union and complementation to which we assign probability), and μ is a *probability measure*, that is, a function $\mu : \mathcal{F} \to [0, 1]$ satisfying the following properties:

P1. $\mu(\emptyset) = 0$.

P2. $\mu(W) = 1$.

P3. $\mu(A \cup B) = \mu(A) + \mu(B)$, if A and B are disjoint elements of $\mathcal{F}$.

If $\mathcal{F} = 2^W$, then it suffices to define a probability measure μ only on the elements of W; it can then be uniquely extended to all subsets of W by taking $\mu(A) = \sum_{w \in A} \mu(w)$.[2]

Next, let us briefly justify probability. The classical approach to applying probability, which goes back to the seventeenth and eighteenth centuries, is to reduce a situation to a number of elementary outcomes.

Definition 1.6. The *principle of indifference* is an assumption that all elementary outcomes are equally likely.

[2] If the argument is a singleton subset $\{w\}$, we often abuse notation and write $\mu(w)$ rather than $\mu(\{w\})$.

Intuitively, in the absence of any other information, there is no reason to consider one more likely than another. Applying the principle of indifference, if there are n elementary outcomes, the probability of each one is $1/n$; the probability of a set of k outcomes is k/n. Clearly, this definition satisfies P2 and P3 (where W consists of all the elementary outcomes).

While taking possible worlds to be equally probable is a very compelling intuition, the trouble with the principle of indifference is that it is not always obvious how to reduce a situation to elementary outcomes that seem equally likely.

What exactly is meant by the elementary outcomes? How should we choose them? And why should they be equally likely? In light of these issues, philosophers and probabilists have tried to find ways of viewing probability that do not depend on assigning elementary outcomes equal likelihood. Perhaps the two most common views are that (1) the numbers represent relative frequencies, and (2) the numbers reflect subjective assessments of likelihood.

Definition 1.7. Let n be a fixed natural number. *Relative-frequency interpretation* is an assumption that if an experiment is repeated n times, then the probability of an event is taken to be the fraction of the n times the event occurred.

It is easy to see that the relative-frequency interpretation of probability satisfies the additive property P3. Moreover, it is closely related to the intuition behind the principle of indifference. For example, in the case of a (fair) coin, roughly half of the outcomes should be heads and half should be tails.

1.3 Ranking Functions

Another approach to representing uncertainty is given by so called (ordinal) ranking functions.

Definition 1.8. A *ranking function (or κ-ranking)* κ on a set W of possible worlds is a function $2^W \to \mathbb{N}^*$, where $\mathbb{N}^* = \mathbb{N} \cup \{\infty\}$, satisfying the following properties:[3]

[3]Based on ideas that go back to [Spo88].

Rk1. $\kappa(\emptyset) = \infty$.

Rk2. $\kappa(W) = 0$.

Rk3. $\kappa(A \cup B) = \min(\kappa(A), \kappa(B))$, if A and B are disjoint.

As with probability, a ranking function is characterised by its behaviour on singletons in finite spaces; $\kappa(A) = \min_{w \in A} \kappa(w)$. To ensure that Rk2 holds, it must be the case that $\min_{w \in W} \kappa(w) = 0$, that is, at least one element in W must have a rank of 0.

Intuitively, a ranking function assigns a degree of surprise to each subset of worlds in W, where 0 means unsurprising and higher numbers denote greater surprise (with ∞ denoting 'so surprising as to be impossible').

1.4 Possibility Measures

Possibility measures are yet another approach to assigning numbers to sets. They are based on ideas of fuzzy logic.

Definition 1.9. Let the set of possible worlds W be finite and all sets measurable. A *possibility measure* $Poss$ on W is a function $2^W \to [0, 1]$ satisfying the following properties:

Poss1. $Poss(\emptyset) = 0$.

Poss2. $Poss(W) = 1$.

Poss3. $Poss(A \cup B) = \max(Poss(A), Poss(B))$, if A and B are disjoint.

The dual of possibility, called *necessity*, is defined in the obvious way:

$$Nec(A) = 1 - Poss(\overline{A}),$$

where $\overline{A}$ is the complement of A in W.

It follows that, like the cases before, if W is finite and all sets are measurable, then a possibility measure can be characterised by its behaviour on singleton sets: $Poss(A) = \max_{w \in A} Poss(w)$. For Poss2 to be true, it must be the case that $\max_{w \in W} Poss(w) = 1$, that is, at least one element in W must have maximum possibility. Moreover, in infinite spaces, in Poss3 supremum is considered instead of maximum.

Perhaps the most common interpretation given to possibility and necessity is that they capture, not a degree of likelihood, but a (subjective) degree of uncertainty regarding the truth of a statement. $Poss(A)$ estimates the degree an agent believes the true world can be in A, while $Nec(A)$ estimates the degree the agent believes the true world should be necessarily in A.

1.5 Dempster-Shafer Belief Functions

The Dempster-Shafer theory of evidence, originally introduced by Arthur Dempster and then developed by Glenn Shafer, provides another approach to attaching likelihood to events. This approach starts out with a belief function (sometimes called a support function).

Definition 1.10. Let all subsets of a set of possible worlds W be measurable. A *belief function* Bel on W is a function $2^W \to [0,1]$ satisfying the following properties:

B1. $Bel(\emptyset) = 0$.

B2. $Bel(W) = 1$.

B3. $Bel(\cup_{i=1}^{n} A_i) \geq \sum_{i=1}^{n} \sum_{\{I \subseteq \{1,\dots,n\}:|I|=i\}} (-1)^{i+1} Bel(\cap_{j \in I} A_j)$, for $n = 1, 2, 3, \dots$.

Unlike the previous three cases, a belief function defined on 2^W cannot be characterised by its behaviour on singleton sets. Thus, its domain must be viewed as being 2^W (or some algebra over W).

1.6 Other Notions of Uncertainty

A preference ordering on W is a strict partial order $\prec$ over W. Intuitively, $w \prec w'$ holds if w is preferred to w'. Preference orders have been used to provide semantics for default (i.e., conditional) statements.

A parameterised probability distribution (PPD) on W is a sequence $\{Pr_i : i \geq 0\}$

of probability measures over W. Such structures provide semantics for defaults in ϵ-semantics.

More about these frameworks can be found in [FrH01].

1.7 Generic Framework: Plausibility Measures

Here we consider an approach that is a generalisation of all the approaches mentioned so far. The basic idea behind plausibility measures is straightforward. A probability measure maps sets in an algebra $\mathcal{F}$ over a set W of possible worlds to $[0,1]$. A plausibility measure is more general; it maps sets in $\mathcal{F}$ to some arbitrary partially ordered set. If Pl is a plausibility measure, $Pl(A)$ denotes the plausibility of A. If $Pl(A) \leq Pl(B)$, then B is at least as plausible as A. Because the ordering is partial, it could be that the plausibility of two different sets is incomparable. An agent may not be prepared to order two sets in terms of plausibility.

Definition 1.11. A *plausibility space* is a tuple $S = (W, \mathcal{F}, Pl)$, where W is a non-empty set of possible worlds, $\mathcal{F}$ is an algebra of measurable subsets of W, and Pl is a *plausibility measure*, that is, a function $Pl : \mathcal{F} \to D$, where D is some domain of plausibility values partially ordered by a relation $\leq_D$ (so that $\leq_D$ is reflexive, transitive, and anti-symmetric). We assume that D is pointed, that is, it contains two special elements $\top_D$ and $\bot_D$, such that $\bot_D \leq d \leq \top_D$ for all $d \in D$. As usual, we define the ordering $<$ by taking $d_1 < d_2$ if and only if $d_1 \leq d_2$ and $d_1 \neq d_2$. The following properties are satisfied:

Pl1. $Pl(\emptyset) = \bot_D$.

Pl2. $Pl(W) = \top_D$.

Pl3. (A1.) If $A \subseteq B$, then $Pl(A) \leq Pl(B)$.

The last property says that a set must be at least as plausible as any of its subsets, that is, plausibility respects subsets.

The domain D is deliberately suppressed from the tuple S, since the choice of D is not significant here. All that matters is the ordering induced by $\leq_D$ on the subsets in $\mathcal{F}$. We may also omit the $\mathcal{F}$ when its role is not that significant, and just denote a plausibility space as a pair (W, Pl) rather than $(W, \mathcal{F}, Pl)$.[4]

Definition 1.12. We say $PL = (W, Pl, \pi)$ is a *plausibility structure*[5], if (W, Pl) is a plausibility space, and $\pi : \Phi \rightarrow 2^W$ is a valuation function that for every $p \in \Phi$ yields the set $\pi(p) \subseteq W$ of states in which p is true.

Clearly probability measures, possibility and necessity measures, and Dempster-Shafer belief functions are all instances of plausibility measures, where we have $D = [0, 1]$, $\bot = 0$, $\top = 1$, and $\leq_D$ is the standard ordering on the reals. Ranking functions are also instances of plausibility measures; in this case, $D = \mathbb{N}^*$ (where $\mathbb{N}^* = \mathbb{N} \cup \infty$), $\bot = \infty$, $\top = 0$, and $x \leq_{\mathbb{N}^*} y$ if and only if $y \leq x$ under the standard ordering on the ordinals.

Moreover, Friedman and Halpern in [FrH01] also show how to map preference orderings and parameterised probability distributions on W to plausibility measures on W in a way that preserves the truth values of defaults.

Plausibility measures are very general. Pl1–3 (see Definition 1.11) are quite minimal requirements, by design, and arguably are the smallest set of properties that a representation of likelihood should satisfy. It is, of course, possible to add more properties, some of which seem quite natural, but these are typically properties that some representation of uncertainty does not satisfy.

What is the advantage of having this generality? For one thing, by using plausibility measures, it is possible to prove general results about properties of representations of uncertainty. That is, it is possible to show that all representations of uncertainty that have property X also have property Y. Since it may be clear that, say, possibility measures and ranking functions have property X, then it immediately follows that both have property Y; moreover, if Dempster-Shafer belief functions do not have property X, the proof may well give a deeper understanding as to why belief functions do not have property Y.

[4] Recall that in cases of ranking functions and possibility measures we usually take $\mathcal{F} = 2^W$.

[5] Note that we use the notion of a structure in the same meaning as a model. Unification of these two notions might have confused the terminology of the individual frameworks.

1.7.1 Qualitative Plausibility Space

Definition 1.13. A plausibility space (W, Pl) is *qualitative* if it also satisfies the following conditions:

A2. If A, B and C are pairwise disjoint sets, $Pl(A \cup B) > Pl(C)$, and $Pl(A \cup C) > Pl(B)$, then $Pl(A) > Pl(B \cup C)$.

A3. If $Pl(A) = Pl(B) = \bot$, then $Pl(A \cup B) = \bot$.

In that case, Pl is called a *qualitative plausibility measure*. We say $PL = (W, Pl, \pi)$ is a *qualitative plausibility structure*, if (W, Pl) is a qualitative plausibility space, and $\pi : \Phi \to 2^W$ is a valuation function that for every $p \in \Phi$ yields the set $\pi(p) \subseteq W$ of states in which p is true. We denote S^{QPL} the class consisting of all qualitative plausibility structures.

Let us further denote S^{PL} the class of all plausibility structures, and S^{Poss}, S^{κ}, S^{p}, S^{ϵ}, the classes that arise from mapping possibility structures, ranking structures, preferential structures, and parameterised probability distribution, respectively, into plausibility structures. Interestingly enough, in [FrH01] Friedman and Halpern have proved that all of them are composed of qualitative plausibility structures.

Theorem 1.14. *Each of* $S^{Poss}, S^{\kappa}, S^{p}, S^{r}$*, and* S^{ϵ} *is a subset of* S^{QPL}*.*

Further we introduce a rather important notion of a set of most plausible worlds. As we shall see in the following chapters, this concept is closely related to the satisfaction condition for (conditional) beliefs. The following proposition shows that there cannot be more than one such set ([FrH97, p. 275]).

Definition 1.15. Let $S = (W, Pl)$ be a qualitative plausibility space. We say that $A \subseteq W$ is a *set of most plausible worlds* if $Pl(A) > Pl(\overline{A})$ (where $\overline{A}$ is the complement of A in W) and for all $B \subset A, Pl(B) \not> Pl(\overline{B})$. That is A is a minimal set of worlds that is more plausible than its complement.

Proposition 1.16. *Let* $S = (W, Pl)$ *be a qualitative plausibility space. If there is a set of most plausible worlds* $A \subseteq W$*, then it must be unique.*

Proof. Suppose that A and A' are both most plausible sets. We now show that $Pl(A \cap A') > Pl(\overline{A \cap A'})$. Since A and A' are both most plausible sets of worlds, this will show that we must have $A = A'$. To see that $Pl(A \cap A') > Pl(\overline{A \cap A'})$, first note that $A \cap A'$, $A - A'$ and $\overline{A}$ are pairwise disjoint. Since A and A' are most plausible sets of worlds, we have that $Pl((A \cap A') \cup (A - A')) = Pl(A) > Pl(\overline{A})$ and $Pl((A \cap A') \cup \overline{A}) \geq Pl((A \cap A') \cup (A' - A)) = Pl(A') > Pl(\overline{A'}) \geq Pl(A - A')$. We can apply A2 to get that $Pl(A \cap A') > Pl((A - A') \cup \overline{A}) = Pl(\overline{A \cap A'})$. □

In finite plausibility spaces (that is, ones with only finitely many worlds), it is easy to see that there is always a (unique) set of most plausible worlds. In general, however, a set of most plausible worlds does not necessarily exist. What property needs to be added to a qualitative plausibility space to ensure that there is a set of most plausible worlds? As it turns out, the sufficient and necessary condition is for the set of possible worlds W to be (converse) well-founded, where the ordering on worlds is induced by their plausibilities.

Proposition 1.17. *Let $S = (W, Pl)$ be a qualitative plausibility space. There is a set of most plausible worlds $A \subseteq W$ if and only if the ordering on worlds W induced by their plausibilities is (converse) well-founded (i.e., there is no infinite ascending chain $Pl(w_0) < Pl(w_1) < Pl(w_2) < \ldots$[6]).*

Proof. The 'if-direction' is straightforward. Since W is (converse) well-founded, just consider a set $A = \{w : \text{there is no } w' \in W \text{ such that } Pl(w') > Pl(w)\}$, which is literally a set of most plausible worlds. The rest follows easily. The opposite 'only if-direction', however, demands a little bit more work.

For contradiction let us consider the space $S_0 = (W, Pl)$ with W containing an infinite number of worlds $W' = \{w_i \mid i \geq 0\}$ ordered in an infinite chain such that $Pl(w_0) < Pl(w_1) < Pl(w_2) < \ldots$. By assumption, we have a set of most plausible worlds A, that is $Pl(A) > Pl(\overline{A})$. It is easy to see that A must contain an infinite number of the worlds from W'. Let us take an arbitrary $w_i \in A \cap W'$. We want to show that $Pl(A - \{w_i\}) > Pl(\overline{A - \{w_i\}})$, because that would be in contradiction with the assumption that A is the set of most plausible worlds. First note that $A - \{w_i\}, \{w_i\}$ and $\overline{A}$ are pairwise disjoint

[6] To be precise, an ordering is well-founded if and only if it contains no countable infinite descending chains; and therefore it may be more natural to consider a preference ordering: $\ldots w_2 \prec w_1 \prec w_0$ as a counterpart to the ordering on plausibilities given above.

sets. The desired inequality immediately follows from $Pl((A - \{w_i\}) \cup \{w_i\}) = Pl(A) > Pl(\overline{A})$ and $Pl((A - \{w_i\}) \cup \overline{A}) = Pl(W - \{w_i\}) > Pl(w_i)$. The former is given by assumption and for the latter there exists $w_{i+1} \in W'$ such that $Pl(W - \{w_i\}) \geq Pl(w_{i+1}) > Pl(w_i)$ by A1. Hence, applying A2 we get that $Pl(A - \{w_i\}) > Pl(\overline{A} \cup \{w_i\}) = Pl(\overline{A - \{w_i\}})$. □

Since we would like there to be a set of most plausible worlds, let us investigate the class of well-founded structures further. What is its relation to the following property of richness?

Definition 1.18. We say that S is *rich* if for every collection $\varphi_1, \ldots \varphi_n,\ n > 1$, of pairwise mutually exclusive and satisfiable propositional formulas, there is a plausibility structure $PL = (W, Pl, \pi) \in S$ such that:

$$Pl([\![\varphi_1]\!]) > Pl([\![\varphi_2]\!]) > \cdots > Pl([\![\varphi_n]\!]) = \bot,$$

where $[\![\varphi]\!] := \{w \in W \mid PL, w \models \varphi\}$.

The richness requirement is quite mild. It says that S does not place a priori constraints on the relative plausibilities of a collection of disjoint sets. Therefore, it should not be surprising that all the classes of plausibility structures mentioned above are indeed rich (proved in [FrH01, p. 13]).

Theorem 1.19. *Each of $S^{Poss}, S^{\kappa}, S^{p}, S^{r}, S^{\epsilon}$, and S^{QPL} is rich.*

Obviously, richness might be too mild to ensure that there is a set of most plausible worlds.

Proposition 1.20. *The class of all well-founded plausibility structures is rich. However, not all rich classes consist of well-founded plausibility structures.*

Proof. First we show that well-foundedness induces richness. Notice that well-foundedness does not restrict the relative plausibilities of a (finite) collection of disjoint sets. Moreover, in order to satisfy richness only a finite number of worlds is required, and thus there is always a well-founded plausibility structure which suffices the particular collection.

On the other hand, richness does not induce well-foundedness. S^p is the class of all preferential structures, thus including also the ones which are not well-founded. According to the previous theorem S^p is rich. Therefore, the result follows. □

It is worth noticing that such non-well-founded structures can come quite handy. Suppose $S = (W, Pl)$ is a qualitative plausibility space with an infinite number of worlds $W = \{w_i : i \geq 1\}$ ordered in an infinite chain such that $Pl(w_1) < Pl(w_2) < Pl(w_3) < \dots$. Let $\varphi_1, \dots, \varphi_n,\ n > 1$, be arbitrary pairwise mutually exclusive and satisfiable propositional formulas. We can construct a mapping π that maps the first $n-1$ worlds in W to a truth assignment such that $[\![\varphi_i]\!] = w_{n-i}$ for all $1 \leq i < n$, and $[\![\varphi_n]\!] = \emptyset$. As we can see this non-well-founded qualitative plausibility space can 'accommodate' (via appropriate mapping) all such collections for an arbitrary $n > 1$.

Well-foundedness might not even be needed under a stronger yet for our purposes suitable assumption: In general, we will assume that the set W of possible worlds is finite.

With intention to focus on ranking functions and possibility measures we can add even more properties. Notice that both of these measures are realised on a totally ordered set ($\mathbb{N}^*$, and $[0, 1]$, respectively) and plausibility of a union is determined by the minimum, and the maximum, respectively, of plausibilities of its components.

Definition 1.21. We say that a plausibility measure Pl is a *ranking* if it satisfies the following two properties:

A4. $\leq_D$ is a total order; that is, either $Pl(A) \leq_D Pl(B)$ or $Pl(B) \leq_D Pl(A)$ for all sets $A, B \in \mathcal{F}$.

A5. $Pl(A \cup B) = \max(Pl(A), Pl(B))$ for all sets $A, B \in \mathcal{F}$.

Notice that ranked plausibility spaces, i.e., where the plausibility measure is a ranking, are indeed qualitative. It is easy to verify that ranking functions and possibility measures are both rankings. These properties shall prove very useful in Chapter 3.

In this chapter we have introduced several 'static' frameworks with plausibility spaces generalising them all. Now let us proceed further and show how we can update uncertainty on each of these frameworks.

Chapter 2

Updating Uncertainty

Agents continually obtain new information and then must update their beliefs to take this new information into account. How this should be done obviously depends in part on how uncertainty is represented. Each of the methods of representing uncertainty considered in Chapter 1 has an associated method for updating.[1]

2.1 Updating Knowledge

We start by examining perhaps the simplest setting, where an agent's uncertainty is captured by a set W of possible worlds, with no further structure. We assume that the agent obtains the information that the actual world is in some subset U of W. The obvious thing to do in that case is to take the set of possible worlds to be $W \cap U$.

Even in this simple setting, three implicit assumptions are worth bringing out:

1. An agent does not forget.
2. What an agent is told is true.
3. The way an agent obtains the new information does not itself give the agent information.[2]

The first assumption can be explained as follows. Suppose at some point the agent has been told that the actual world is in $U_1, \dots, U_n$. The agent should then consider possible precisely the worlds in $U_1 \cap \cdots \cap U_n$. If she is then told V, she considers possible $U_1 \cap \cdots \cap U_n \cap V$. This seems to justify the idea of capturing

[1] The following overview is taken from [Hal03, ch. 3].

[2] Indeed, there are exceptions to all three of them, e.g., memory-free agents, lying to an agent, and public announcement made by an insider (that is, by one of the agents, not an outsider who sees the whole situation 'from above'), respectively.

updating by U as intersecting the current set of possible worlds with U. (It is not so clear that intersection is appropriate if forgetting is allowed.)

The second assumption is perhaps more obvious but nonetheless worth stressing. What it says is that if the agent is told U then the actual world is in U. We also assume that the agent initially considers the actual world possible. From this it follows that if U_0 is the agent's initial set of possible worlds and she is told U, then $U_0 \cap U \neq \emptyset$ (since the actual world is in $U_0 \cap U$). (It is not even clear how to interpret a situation where the agent's set of possible worlds is empty. If the agent can be told inconsistent information, then clearly intersection is simply not an appropriate way of updating.)

The third assumption simply regulates that making an observation may give more information than just the fact that what is observed is true. If this is not taken into account, intersecting may give an inappropriate answer. (e.g., Finding out that my friend has my book may also give me an extra piece of information that she has stopped by at my place where I left it.)

2.2 Probabilistic Conditioning

Suppose that an agent's uncertainty is replaced by a probability measure μ on W and then the agent observes or learns (that the actual world is in) U. How should μ be updated to a new probability measure $\mu|U$ that takes this new information into account? Clearly if the agent believes that U is true, then it seems reasonable to require that all the worlds in $\overline{U}$ are impossible:

$$\mu(\overline{U} \mid U) = 0, \tag{2.1}$$

where we write $\mu(\overline{U} \mid U)$ rather than $\mu|U(\overline{U})$.

What about worlds in U? What should their probability be? One reasonable intuition about the worlds in U is that the relative likelihood of worlds in U should remain unchanged. That is, if $V_1, V_2 \subseteq U$ with $\mu(V_2) > 0$, then

$$\frac{\mu(V_1)}{\mu(V_2)} = \frac{\mu(V_1 \mid U)}{\mu(V_2 \mid U)}. \tag{2.2}$$

As proved in [Hal03, p. 72], the equations (2.1) and (2.2) completely determine $\mu|U$ if $\mu(U) > 0$.

Proposition 2.1. *If* $\mu(U) > 0$ *and* $\mu|U$ *is a probability measure on* W *satisfying* (2.1) *and* (2.2), *then*

$$\mu(V\ |U) = \frac{\mu(V \cap U)}{\mu(U)}. \tag{2.3}$$

Here $\mu|U$ is called a conditional probability (measure), and $\mu(V\ |U)$ is read 'the probability of V given (or conditional on) U'. Sometimes $\mu(U)$ is called the unconditional probability of U.

Conditioning is a wonderful tool, but it does suffer from some problems, particularly when it comes to dealing with events with probability 0. Traditionally, (2.3) is taken as the definition of $\mu(V\ |U)$ if μ is an unconditional probability measure and $\mu(U) > 0$; if $\mu(U) = 0$, then the conditional probability $\mu(V\ |U)$ is undefined. This leads to a number of philosophical difficulties regarding worlds (and sets) with probability 0. Are they really impossible? If not, how unlikely does a world have to be before it is assigned probability 0? Should a world ever be assigned probability 0? If there are worlds with probability 0 that are not truly impossible, then what does it mean to condition on sets with probability 0?

Some of these issues can be sidestepped by treating conditional probability, not unconditional probability, as the basic notion. A conditional probability measure takes pairs U, V of subsets as arguments; $\mu(V,\ U)$ is generally written as $\mu(V\ |U)$ to stress the conditioning aspects. What pairs $(V,\ U)$ should be allowed as arguments to μ? The intuition is that for each fixed second argument U, the function $\mu(.\ ,U)$ should be a probability measure. In order to satisfy this requirement we make use of the following definition.

Definition 2.2. A *Popper algebra* over W is a set $\mathcal{F} \times \mathcal{F}'$ of subsets of $W \times W$ such that

Acc1. $\mathcal{F}$ is an algebra over W.

Acc2. $\mathcal{F}'$ is a nonempty subset of $\mathcal{F}$.

Acc3. $\mathcal{F}'$ is closed under supersets in $\mathcal{F}$: if $V \in \mathcal{F}', V \subseteq V'$, and $V' \in \mathcal{F}$ then $V' \in \mathcal{F}'$.

Notice that $\mathcal{F}'$ need not be an algebra.

Definition 2.3. A *conditional probability space* is a tuple $(W, \mathcal{F}, \mathcal{F}', \mu)$ such that $\mathcal{F} \times \mathcal{F}'$ is a Popper algebra over W and $\mu : \mathcal{F} \times \mathcal{F}' \to [0, 1]$ is *conditional probability measure* that satisfies the following conditions:

CP1. $\mu(U \mid U) = 1$ if $U \in \mathcal{F}'$.

CP2. $\mu(V_1 \cup V_2 \mid U) = \mu(V_1 \mid U) + \mu(V_2 \mid U)$ if $V_1 \cap V_2 = \emptyset,\ V_1,\ V_2 \in \mathcal{F}$ and $U \in \mathcal{F}'$.

CP3. $\mu(U_1 \cap U_2 \mid U_3) = \mu(U_1 \mid U_2 \cap U_3) \times \mu(U_2 \mid U_3)$ if $U_1 \in \mathcal{F}$ and $U_2 \cap U_3 \in \mathcal{F}'$.[3]

CP4. $\mu(V \mid U) = \mu(V \cap U \mid U)$ if $U \in \mathcal{F}'$ and $V \in \mathcal{F}$.

CP5. $\mu(U_1 \mid U_3) = \mu(U_1 \mid U_2) \times \mu(U_2 \mid U_3)$, if $U_1 \subseteq U_2 \subseteq U_3,\ U_2,\ U_3 \in \mathcal{F}'$ and $U_1 \in \mathcal{F}$.

CP1 and CP2 are just the obvious analogues of P2 and P3 (see Definition 1.5). CP3 is perhaps best understood through CP4 and CP5, since in the presence of CP1, CP3 is equivalent to CP4 and CP5. CP4 just says that, when conditioning on U, everything should be relativised to U. CP5 says that if $U_1 \subseteq U_2 \subseteq U_3$, it is possible to compute the conditional probability of U_1 given U_3 by computing the conditional probability of U_1 given U_2, computing the conditional probability of U_2 given U_3, and then multiplying them together.

If μ is a conditional probability measure, then we usually write $\mu(U)$ instead of $\mu(U \mid W)$. Thus, in the obvious way, a conditional probability measures determines an unconditional probability measure.

Probabilistic conditioning can be justified in much the same way that probability is justified. For example, if it seems reasonable to apply the principle of indifference (i.e., a natural assumption that all elementary outcomes are equally likely) to W and then U is observed or learned, it seems equally reasonable to apply the principle of indifference again on $W \cap U$. This results in taking all the elements of $W \cap U$ to be equally likely and assigning all the elements in $\overline{W \cap U}$ probability 0, which is exactly what (2.3) says. Similarly, using the relative-frequency interpretation, $\mu(V \mid U)$ can be viewed as the fraction of times that V occurs of the times that U occurs. Again, (2.3) holds.

[3]Notice that here $\times$ is used to denote multiplying; not to be confused with the cases of a Cartesian product.

One of the most important and widely applicable results in probability theory is called Bayes' Rule. It relates $\mu(V \mid U)$ and $\mu(U \mid V)$.[4]

Proposition 2.4 (Bayes' Rule)**.** *Given a conditional probability measure μ on W and $\mu(U)$, $\mu(V) > 0$, the following holds:*

$$\mu(V \mid U) = \frac{\mu(U \mid V)\mu(V)}{\mu(U)}.$$

2.3 Conditioning Ranking Functions

Defining conditional κ-rankings is straightforward, using an analogue of the properties CP1–3 that were used to characterise probabilistic conditioning above.

Definition 2.5. A *conditional ranking function* κ is a function mapping a Popper algebra $2^W \times \mathcal{F}' \to \mathbb{N}^*$ satisfying the following properties:

CRk1. $\kappa(\emptyset \mid U) = \infty$ if $U \in \mathcal{F}'$.

CRk2. $\kappa(U \mid U) = 0$ if $U \in \mathcal{F}'$.

CRk3. $\kappa(V_1 \cup V_2 \mid U) = \min(\kappa(V_1 \mid U), \kappa(V_2 \mid U))$ if $V_1 \cap V_2 = \emptyset, V_1, V_2 \in \mathcal{F}$ and $U \in \mathcal{F}'$.

CRk4. $\kappa(U_1 \cap U_2 \mid U_3) = \kappa(U_1 \mid U_2 \cap U_3) + \kappa(U_2 \mid U_3)$ if $U_1 \in \mathcal{F}$ and $U_2 \cap U_3 \in \mathcal{F}'$.

Given an unconditional ranking function κ, the unique conditional ranking function with these properties with domain $2^W \times \mathcal{F}'$, where $\mathcal{F}' = \{U : \kappa(U) \neq \infty\}$ is defined via

$$\kappa(V \mid U) = \kappa(V \cap U) - \kappa(U).$$

This, indeed, is a motivation for choosing $+$ as the replacement for probabilistic $\times$ in CRk4.

Notice that there is an obvious analogue of Bayes' Rule for ranking functions:

$$\kappa(V \mid U) = \kappa(U \mid V) + \kappa(V) - \kappa(U).$$

[4]The following proposition with a straightforward proof can be found in [Hal03, p. 79].

2.4 Conditioning Possibility Measures

The definition of conditional possibility we present here takes as its point of departure the fact that minimum should play the same role in the context of possibility as multiplication does for probability (and addition for ranking functions). In the case of probability, this role is characterised by CP3.

Definition 2.6. A *conditional possibility measure* $Poss$ is a function mapping a Popper algebra $2^W \times \mathcal{F}' \to [0,1]$ satisfying the following properties:

CPoss1. $Poss(\emptyset \mid U) = 0$ if $U \in \mathcal{F}'$.

CPoss2. $Poss(U \mid U) = 1$ if $U \in \mathcal{F}'$.

CPoss3. $Poss(V_1 \cup V_2 \mid U) = \max(Poss(V_1 \mid U), Poss(V_2 \mid U))$ if $V_1 \cap V_2 = \emptyset$, $V_1, V_2 \in \mathcal{F}$ and $U \in \mathcal{F}'$.

CPoss4. $Poss(U_1 \cap U_2 \mid U_3) = \min(Poss(U_1 \mid U_2 \cap U_3), Poss(U_2 \mid U_3))$ if $U_1 \in \mathcal{F}$ and $U_2 \cap U_3 \in \mathcal{F}'$.

CPoss4 is just the result of replacing μ by $Poss$ and $\times$ by min in CP3 (cf. Definition 2.3).

One approach that has been taken in the literature to defining a canonical conditional possibility measure determined by an unconditional possibility measure is to make things 'as possible as possible.' That is, given an unconditional possibility measure $Poss$, the largest conditional possibility measure $Poss(. \mid .)$ consistent with CPoss1–4 that is an extension of $Poss$ is considered. This leads to the following definition in the case that $Poss(U) > 0$:

$$Poss(V \mid U) = \begin{cases} Poss(V \cap U) & \text{if } Poss(V \cap U) < Poss(U) \\ 1 & \text{if } Poss(V \cap U) = Poss(U). \end{cases}$$

With this definition, there is no direct analogue to Bayes' Rule; $Poss(V \mid U)$ is not determined by $Poss(U \mid V), Poss(U)$, and $Poss(V)$. However, it is immediate from CPoss4 that there is still a close relationship among $Poss(V \mid U)$, $Poss(U \mid V)$, $Poss(U)$, and $Poss(V)$ that is somewhat akin to Bayes' Rule, namely,

$$\min(Poss(V \mid U), Poss(U)) = \min(Poss(U \mid V), Poss(V)) = Poss(V \cap U).$$

2.5 Conditioning Belief Functions and Others

We have omitted the conditioning method for belief functions, since it requires a little bit more technical details which are beyond the scope of this overview. The method can be found in [Hal03, ch. 3] together with those for conditioning with sets of probabilities, and conditioning inner and outer measures.

2.6 Generic Framework: Conditional Plausibility Space

We proceed in a manner similar in spirit to that for probability.

Definition 2.7. A *conditional plausibility space* (cps) is a tuple $(W, \mathcal{F}, \mathcal{F}', Pl)$, where $\mathcal{F} \times \mathcal{F}'$ is a Popper algebra over W, D is a partially ordered set of plausibility values, and $Pl : \mathcal{F} \times \mathcal{F}' \to D$ is a *conditional plausibility measure* (cpm) that satisfies the following conditions:

CPl1. $Pl(\emptyset \mid U) = \bot$.

CPl2. $Pl(U \mid U) = \top$.

CPl3. If $V \subseteq V'$, then $Pl(V \mid U) \leq Pl(V' \mid U)$.

CPl4. $Pl(V \mid U) = Pl(V \cap U \mid U)$.

CPl1–3 just say that Pl1–3 (see Definition 1.11) hold for $Pl(. \mid U)$, so that $Pl(. \mid U)$ is a plausibility measure for each fixed $U \in \mathcal{F}'$. CPl4 is the obvious analogue to CP4 (see Definition 2.3). Since there is no notion of multiplication for plausibility measures yet, it is not possible to give an analogue of CP3 for conditional plausibility.

Definition 2.8. A cps $(W, \mathcal{F}, \mathcal{F}', Pl)$ is *acceptable* if $V \in \mathcal{F}$, $U \in \mathcal{F}'$, and $Pl(V \mid U) \neq \bot$ implies $V \cap U \in \mathcal{F}'$.

Acceptability is a generalisation of the observation that if $\mu(V) \neq 0$, then conditioning on V should be defined. It says that if $Pl(V \mid U) \neq \bot$, then conditioning on $V \cap U$ should be defined. All the constructions that were used for defining conditional likelihood measures result in acceptable cps's. On the other hand, acceptability is not required in the definition of conditional probability space (cf. Definition 2.3).

CPl1–4 are rather minimal requirements. Should there be others? The following coherence condition, which relates conditioning on two different sets, seems quite natural:

CPl5. If $U \cap U' \in \mathcal{F}'$, $U, U', V, V' \in \mathcal{F}$, and $Pl(U \mid U') \neq \bot$, then $Pl(V \mid U \cap U') \leq Pl(V' \mid U \cap U')$ iff $Pl(V \cap U \mid U') \leq Pl(V' \cap U \mid U')$.

It can be shown that CPl5 implies CPl4 if $\mathcal{F}'$ has the property that characterises acceptable cps's. While CPl5 seems quite natural, and it holds for all conditional probability measures, conditional possibility measures, and conditional ranking functions, it does not hold in general for belief functions.

Given an unconditional plausibility measure Pl, is it possible to construct a conditional plausibility measure extending Pl? It turns out that it is. The idea is quite straightforward. Given an unconditional plausibility measure Pl defined on an algebra $\mathcal{F}$ with range D, for each set $U \in \mathcal{F}$, start by defining a new plausibility measure Pl_U with range $D_U = \{d \in D : d \leq Pl(U)\}$ by taking $Pl_U(V) = Pl(V \cap U)$. Note that $\top_{D_U} = Pl(U)$. Thus, defining $Pl(V \mid U)$ as $Pl_U(V)$ will not quite work, because then CPl2 is not satisfied; in general, $Pl_U(W) \neq Pl_V(W)$. In [Hal03, pp. 99–101] it is shown how to overcome this inconvenience and get the desired cps.

2.6.1 Algebraic Conditional Plausibility Space

The definitions of conditional ranking and conditional possibility were motivated in part by considering analogues for ranking and possibility of addition and multiplication in probability. In general, many plausibility spaces of interest have 'more structure'. There are analogues of addition and multiplication which could be added to plausibility.

Definition 2.9. A cps $(W, \mathcal{F}, \mathcal{F}', Pl)$ where Pl has range D is *algebraic* if it is acceptable and there are functions $\oplus$ and $\otimes$ mapping $D \times D \to D$ such that the following properties hold:

Alg1. Pl is additive with respect to $\oplus$, that is, $Pl(V_1 \cup V_2 \mid U) = Pl(V_1 \mid U) \oplus Pl(V_2 \mid U)$ if $V_1 \cap V_2 = \emptyset$, $V_1, V_2 \in \mathcal{F}$ and $U \in \mathcal{F}'$.

Alg2. $Pl(U_1 \cap U_2 \mid U_3) = Pl(U_1 \mid U_2 \cap U_3) \otimes Pl(U_2 \mid U_3)$ if $U_1 \in \mathcal{F},\ U_2 \cap U_3 \in \mathcal{F}'$.

Alg3. $\otimes$ distributes over $\oplus$; that is, $a \otimes (b_1 \oplus \cdots \oplus b_n) = (a \otimes b_1) \oplus \cdots \oplus (a \otimes b_n)$ if $(a, b_1), \ldots, (a, b_n), (a, b_1 \oplus \cdots \oplus b_n) \in \mathrm{Dom}(\otimes)$ and $(b_1, \ldots, b_n), (a \otimes b_1, \ldots, a \otimes b_n) \in \mathrm{Dom}(\oplus)$, where $\mathrm{Dom}(\oplus) = \{(Pl(V_1 \mid U), \ldots, Pl(V_n \mid U)) : V_1, \ldots, V_n \in \mathcal{F}$ are pairwise disjoint and $U \in \mathcal{F}'\}$ and $\mathrm{Dom}(\otimes) = \{(Pl(U_1 \mid U_2 \cap U_3), Pl(U_2 \mid U_3)) : U_1 \in \mathcal{F},\ U_2 \cap U_3 \in \mathcal{F}'\}$.

Alg4. If $(a, c), (b, c) \in \mathrm{Dom}(\otimes), a \otimes c \leq b \otimes c$, and $c \neq \bot$, then $a \leq b$.

If $(W, \mathcal{F}, \mathcal{F}', Pl)$ is an algebraic cps, then Pl is called an *algebraic cpm.*

Alg1 and Alg2 are clearly analogues of CP2 and CP3 (cf. Definition 2.3).

Probability measures, ranking functions, and possibility measures are all additive. In the case of probability measures, $\oplus$ is $+$; in the case of ranking functions, it is min; in the case of possibility measures, it is max. However, belief functions are not additive, and thus they cannot be algebraic.[5]

Proposition 2.10. *The constructions for extending an unconditional probability measure, ranking function, and possibility measure to a cps result in algebraic cps's.*

Many of the properties that are associated with (conditional) probability hold more generally for algebraic cps's. We consider three of them here ([Hal03, pp. 103–104]). The first two say that $\bot$ and $\top$ act like 0 and 1 with respect to addition and multiplication. Let $Range(Pl) = \{d : Pl(V \mid U) = d$ for some $(V,\ U) \in \mathcal{F} \times \mathcal{F}'\}$.

Lemma 2.11. *If $(W, \mathcal{F}, \mathcal{F}', Pl)$ is an algebraic cps, then $d \oplus \bot = \bot \oplus d = d$ for all $d \in Range(Pl)$.*

Lemma 2.12. *If $(W, \mathcal{F}, \mathcal{F}', Pl)$ is an algebraic cps, then for all $d \in Range(Pl)$:*

(i) $d \otimes \top = d$.

[5] In [Hal03, p. 102] a richer version of the following proposition can be found together with its proof.

(ii) *if* $d \neq \bot$, *then* $\top \otimes d = d$.

(iii) *if* $d \neq \bot$, *then* $\bot \otimes d = \bot$.

(iv) *if* $(d, \bot) \in \mathrm{Dom}(\otimes)$, *then* $\top \otimes \bot = d \otimes \bot = \bot \otimes \bot = \bot$.

And the third property is an analogue of a standard property for probability that shows how $Pl(V \mid U)$ can be computed by partitioning into subsets.

Lemma 2.13. *Suppose that* $(W, \mathcal{F}, \mathcal{F}', Pl)$ *is an algebraic cps,* $A_1, \ldots, A_n$ *is a partition of* W, $A_1, \ldots, A_n \in \mathcal{F}$, *and* $U \in \mathcal{F}'$. *Then*

$$Pl(V \mid U) = \oplus_{\{i: A_i \cap U \in \mathcal{F}'\}} Pl(V \mid A_i \cap U) \otimes Pl(A_i \mid U).$$

We conclude this section by abstracting a property that holds for all the constructions of cps's from unconditional plausibility measures (i.e., the constructions given in the case of probability, ranking functions, possibility, and plausibility).

Definition 2.14. A cps $(W, \mathcal{F}, \mathcal{F}', Pl)$ is *standard* if $\mathcal{F}' = \{U : Pl(U) \neq \bot\}$.

2.6.2 Jeffrey's Rule

Up to now, we have assumed that the information received is of the form "the actual world is in U". But information does not always come in such nice packages.[6]

Example 2.15. Suppose that an object is either red, blue, green, or yellow. An agent initially ascribes probability 1/5 to each of red, blue, and green, and probability 2/5 to yellow. Then the agent gets a quick glimpse of the object in a dimply lit room. As a result of this glimpse, she believes that the object is probably a darker colour, although she is not sure. She thus ascribes probability .7 to it being green or blue and probability .3 to it being red or yellow. How should she update her initial probability measure based on this observation?

Note that if the agent had definitely observed that the object was either blue or green, she would update her belief by conditioning on $\{blue, green\}$. However,

[6] This topic can be found in [Hal03, pp. 105–107].

her observation was not good enough to confirm that the object was definitely blue or green (nor that it was red or yellow). Rather, it can be represented as $.7\{blue, green\}; .3\{red, yellow\}$. This suggests that an appropriate way of updating the agent's initial probability measure μ is to consider the linear combination $\mu' = .7\mu|\{blue, green\} + .3\mu|\{red, yellow\}$. As expected, $\mu'(\{blue, green\}) = .7$ and $\mu'(\{red, yellow\}) = .3$. Moreover, $\mu'(red) = .1$, $\mu'(yellow) = .2$, and $\mu'(blue) = \mu'(green) = .35$. Thus, μ' gives the two sets about which the agent has information – $\{blue, green\}$ and $\{red, yellow\}$ – the expected probabilities. Within each of these sets, the relative probability of the outcomes remains the same as before conditioning.

More generally, suppose that $U_1, \dots U_n$ is a partition of W (i.e., $\cup_{i=1}^n U_i = W$ and $U_i \cap U_j = \emptyset$ for $i \neq j$) and the agent observes $\alpha_1 U_1; \dots; \alpha_n U_n$, where $\alpha_1 + \dots + \alpha_n = 1$. This is to be interpreted as an observation that leads the agent to believe U_j with probability α_j, for $j = 1, \dots, n$. Moreover, since the observation does not give any extra information regarding subsets of U_j, the relative likelihood of worlds in U_j should remain unchanged. This suggests that $\mu|(\alpha_1 U_1; \dots; \alpha_n U_n)$, the probability measure resulting from the update, should have the following property for $j = 1, \dots, n$:

J. $\mu|(\alpha_1 U_1; \dots; \alpha_n U_n)(V) = \alpha_j \frac{\mu(V)}{\mu(U_j)}$ if $V \subseteq U_j$ and $\mu(U_j) > 0$.

Taking $V = U_j$ in J, it follows that

J1. $\mu|(\alpha_1 U_1; \dots; \alpha_n U_n)(U_j) = \alpha_j$ if $\mu(U_j) > 0$.

Moreover, if $\alpha_j > 0$, the following analogue of (2.2) is a consequence of J (and J1):

J2. $\frac{\mu(V)}{\mu(U_j)} = \frac{\mu|(\alpha_1 U_1; \dots; \alpha_n U_n)(V)}{\mu|(\alpha_1 U_1; \dots; \alpha_n U_n)(U_j)}$ if $V \subseteq U_j$ and $\mu(U_j) > 0$.

Property J uniquely determines what is known as Jeffrey's Rule of conditioning.[7]

Definition 2.16. We define *Jeffrey's Rule* of conditioning as follows:

$$\mu|(\alpha_1 U_1; \dots; \alpha_n U_n)(V) = \alpha_1 \mu(V \mid U_1) + \dots + \alpha_n \mu(V \mid U_n).$$

We take $\alpha_j \mu(V \mid U_j)$ to be 0 here if $\alpha_j = 0$, even if $\mu(U_j) = 0$.

[7] It was defined by Richard Jeffrey.

Definition 2.17. An observation is *consistent* with the initial probability, if it does not give positive probability to a set that was initially thought to have probability 0. Formally, if $\alpha_j > 0$ then $\mu(U_j) > 0$.

As long as the observation is consistent with the initial probability Jeffrey's Rule is defined and gives the unique probability measure satisfying property J.

Note that $\mu|U = \mu|(1U; 0\overline{U})$, so the usual notion of probabilistic conditioning is just a special case of Jeffrey's Rule. However, probabilistic conditioning has one attractive feature that is not maintained in the more general setting of Jeffrey's Rule. Suppose that the agent makes two observations, U_1 and U_2. It is easy to see that if $\mu(U_1 \cap U_2) \neq 0$, then

$$(\mu|U_1)|U_2 = (\mu|U_2)|U_1 = \mu|(U_1 \cap U_2)$$

That is, the following three procedures give the same result: (a) condition on U_1 and then U_2, (b) condition on U_2 and then U_1, and (c) condition on $U_1 \cap U_2$ (which can be viewed as conditioning simultaneously on U_1 and U_2). The analogous result does not hold for Jeffrey's Rule, because according to its definition the last observation determines the probability of a certain set, so the order of observation matters. Thus, if $O_1 \neq O_2$ then $(\mu|O_1)|O_2 \neq (\mu|O_2)|O_1$.

There are straightforward analogues of Jeffrey's Rule for ranking functions and possibility measures.

- For ranking functions, the analogue is based on the observation that $+$ and $\times$ for probability become min and $+$ for ranking, and the role of 1 is played by 0. Thus, for an observation of the form $\alpha_1 U_1; \ldots; \alpha_n U_n$, where $\alpha_i \in \mathbb{N}^*, i = 1, \ldots, n$ and $\min(\alpha_1, \ldots, \alpha_n) = 0$,

$$\kappa|(\alpha_1 U_1; \ldots; \alpha_n U_n)(V) = \min(\alpha_1 + \kappa(V\ |U_1), \ldots, \alpha_n + \kappa(V\ |U_n)).$$

- For possibility measures, $+$ and $\times$ become max and min. Thus, for an observation of the form $\alpha_1 U_1; \ldots; \alpha_n U_n$, where $\alpha_i \in [0, 1]$ for $i = 1, \ldots, n$ and $\max(\alpha_1, \ldots, \alpha_n) = 1$,

$$\begin{aligned} &Poss|(\alpha_1 U_1; \ldots; \alpha_n U_n)(V) \\ &= \max(\min(\alpha_1, Poss(V\ |U_1)), \ldots, \min(\alpha_n, Poss(V\ |U_n))). \end{aligned}$$

In this chapter we have introduced the methods for updating uncertainty on the frameworks from Chapter 1. We have shown that conditional plausibility spaces generalise the conditional extensions of all these frameworks. Based on the ideas from probability, we have also formalised some additional algebraic properties for cps's.

"Conditioning is a wonderful tool, but..." [Hal03, p. 74] As the last subsection shows, we can also update with weaker notions than knowledge when only 'incomplete information' in form of belief is available. In those cases Jeffrey's Rule might be applied.

Chapter 3

Public Announcement by Conditioning

In this main chapter we show how public announcement on three different frameworks can be embedded in a conditional plausibility space. We focus on frameworks based on single-agent plausibility spaces (in [BaS08]), ranking functions, and possibility measures (from Chapter 1), respectively. Each of them is studied in a separate section. First we present their corresponding update mechanisms. Then we list the properties which need to be added to a cps in order to recover the particular framework. At the end of each section we prove that the appropriate cps and the space corresponding to the framework in question remain 'equivalent' after public announcement.

Three qualitative belief revision policies have received substantial attention in dynamic epistemic logic: conditioning, lexicographic revision (also known as radical revision), and minimal revision (also known as conservative revision). In particular, the first may be related (via its eliminative nature) to public announcement.[1]

Definition 3.1. Take a conditional plausibility space $(W, \mathcal{F}, \mathcal{F}', Pl)$ and a set $P \in \mathcal{F}'$ (representing the set of all the worlds where the corresponding proposition P is true). Below we will call any $w \in W \cap P$ a 'p-world'.[2]

- *Conditioning* of the conditional plausibility space $(W, \mathcal{F}, \mathcal{F}', Pl)$ with the set P results in removing all inconsistencies with P, i.e., the operation gives a new conditional plausibility space $(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P)$, where W_P includes only the p-worlds, $\mathcal{F}_P$ and $\mathcal{F}'_P$ include only the sets of the p-worlds, and Pl_P is cut down to the new domain $\mathcal{F}_P \times \mathcal{F}'_P$.

- *Lexicographic revision* of the conditional plausibility space $(W, \mathcal{F}, \mathcal{F}', Pl)$ with the set P results in keeping the same states in W but promoting all the p-worlds to be more plausible than all those that are not p-worlds, and within the two clusters the order remains unchanged.

[1] The idea of the following definition comes from [GHJ14].

[2] For reasons of convenience we use a standard notation p here, but note that a set P can be represented by all the propositional formulas, not just atomic. We do not consider modal formulas to avoid the issue of the well-known 'Moore-sentences'; see [DHK06].

- *Minimal revision* of the conditional plausibility space $(W, \mathcal{F}, \mathcal{F}', Pl)$ with the set P results in promoting the most plausible p-worlds to be the most plausible overall, the rest of the order remaining the same. As in the case of lexicographic revision, W stays the same throughout the process.

In this chapter we deal primarily with conditioning, since the other two types of revision require a certain update model using plausibility measures. However, we illustrate such a general model in Chapter 4.

Next, let us introduce public announcement on a conditional plausibility space, which is basically an update by conditioning with P.

Definition 3.2. *Public announcement* $!P$ is a function R_{PA} that updates a conditional plausibility state, i.e., it associates to any conditional plausibility space $(W, \mathcal{F}, \mathcal{F}', Pl)$ and any set P, such that $W \cap P \in \mathcal{F}'$, some new conditional plausibility space

$$R_{PA}((W, \mathcal{F}, \mathcal{F}', Pl), P) := (W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P),$$

where $W_P = W \cap P$, $\mathcal{F}_P = \{U \cap P : U \in \mathcal{F}\}$, $\mathcal{F}'_P = \{U \cap P : U \in \mathcal{F}' \;\&\; Pl(U \,|P) \neq \bot\}$, and Pl_P is Pl restricted to $\mathcal{F}_P \times \mathcal{F}'_P$, so that $Pl_P(U) = Pl(U \,|P)$ (by CPl4; see Definition 2.7).

It is easy to see that the resulting $(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P)$ is also a conditional plausibility space, thus the update by public announcement is well defined. Notice that in principle we want to condition on P and take $Pl(. \,|P)$, but we also need to leave the rest of the subsets of P in $\mathcal{F}'_P$ in order to be able to apply several plausible public announcements one after another.

Now, let us define what it means for a formula $\varphi \in \mathcal{L}_\Phi$ to be true at a world w in a conditional plausibility structure PL.

Definition 3.3 (Satisfaction relation)**.** Let $PL = (W, 2^W, \mathcal{F}', Pl, \pi)$ be a conditional plausibility structure. The *satisfaction relation* for the propositional connectives is standard (see Definition 1.3), and for the conditional belief is given by, for $w \in W$, $\varphi \in \mathcal{L}_\Phi$,

$$PL, w \models B^\psi \varphi \;\text{ iff }\; Pl([\![\psi]\!]) = \bot \text{ or } Pl([\![\varphi \wedge \psi]\!]) >_D Pl([\![\neg\varphi \wedge \psi]\!]),$$

where $[\![\psi]\!] := \{w \in W : PL, w \models \psi\}$.

In the presence of CPl5 we can equivalently write:

$$PL, w \models B^{\psi}\varphi \text{ iff } Pl([\![\psi]\!]) = \bot \text{ or } Pl([\![\varphi]\!] \mid [\![\psi]\!]) >_D Pl([\![\neg\varphi]\!] \mid [\![\psi]\!]).$$ [3]

The implicit assumption here is that $[\![\psi]\!] \in \mathcal{F}$ iff $Pl([\![\psi]\!]) \neq \bot$.

Recall that "... [c]onditional beliefs 'pre-encode' beliefs that [an agent] would have if [she] learnt certain things" [Ben04, p. 11]. Here the learnt information is denoted by conditioning.

For the obvious similarities with default semantics in [FrH01] let us briefly review the matter here.

3.1 Defaults

Defaults are statements of the form "if ψ then typically/ likely/ by default φ", which is denoted $\psi \to \varphi$. For example, the default "birds typically fly" is represented *Bird* $\to$ *Fly*. According to the semantics stated above this is exactly how we deal with a conditional belief $B^{\psi}\varphi$. It means that $\varphi \wedge \psi$ is more likely to be the case than $\neg\varphi \wedge \psi$ (i.e., if ψ then likely φ).

As it turns out our focus will eventually fall on conditional beliefs. Hence, we state the following definitions for the upcoming purposes.

Definition 3.4. We say that two plausibility spaces (W, Pl) and (W, Pl') (resp. two plausibility structures (W, Pl, π) and (W, Pl', π)) are *order-equivalent* if for any $A, B \subseteq W$, we have $Pl(A) \leq_D Pl(B)$ if and only if $Pl'(A) \leq_D Pl'(B)$.

If all that we are interested in is default reasoning (including our conditional beliefs), then all that matters is the relative plausibility of disjoint sets.

Definition 3.5. We say that two plausibility spaces (W, Pl) and (W, Pl') (resp. two plausibility structures (W, Pl, π) and (W, Pl', π)) are *default-equivalent* if for all disjoint subsets A and B of W, we have $Pl(A) <_D Pl(B)$ if and only if $Pl'(A) <_D Pl'(B)$.[4]

Clearly, if structures (W, Pl, π) and (W, Pl', π) are default-equivalent, then they satisfy the same defaults, hence the same conditional beliefs in our case.

[3] As expected $[\![\varphi]\!] \cap [\![\psi]\!] = [\![\varphi \wedge \psi]\!]$.

[4] $<$ is the strict partial order determined by the corresponding $\leq$.

3.2 Single-Agent Plausibility Space

Here we present a single-agent plausibility space (saps) as defined by Baltag and Smets in [BaS08].[5]

Definition 3.6. A *single-agent plausibility space (or frame)* is a tuple $(W, \leq)$ consisting of a set W of states and a well-preorder $\leq$, i.e., a reflexive, transitive binary relation on W such that every non-empty subset has minimal elements. Using the notation

$$Min_{\leq} P := \{w \in P : w \leq w' \text{ for all } w' \in P\}$$

for the set of $\leq$-minimal elements of P, the last condition says that for every set $P \subseteq W$, if $P \neq \emptyset$ then $Min_{\leq} P \neq \emptyset$.

The usual reading of $w \leq v$ is that state w is at least as plausible as state v. The minimal states in $Min_{\leq} P$ are thus the most plausible states satisfying proposition P. As usually, we write $w < v$ iff $w \leq v$ but $v \not\leq w$, for the *strict plausibility relation*. Similarly, we write $w \cong v$ iff both $w \leq v$ and $v \leq w$, for the *equi-plausibility* (or indifference) relation.

Observe that a single-agent plausibility space is just a special case of a Kripke frame. So, as it is standard for Kripke frames in general, we can define a plausibility model as follows.

Definition 3.7. Given a single-agent plausibility space $(W, \leq)$ a *W-proposition* is any subset $P \subseteq W$. Intuitively, we say that a state w satisfies the proposition P if $w \in P$. Then we can define a *single-agent plausibility model (or structure)* to be a tuple $(W, \leq, \pi)$, consisting of a single-agent plausibility space $(W, \leq)$ together with a valuation function $\pi : \Phi \to 2^W$ mapping every element of a given set Φ of atomic propositions into W-propositions.[6]

Let us give a brief interpretation. The elements of W will represent the possible worlds. The atomic propositions $p \in \Phi$ represent 'ontic' (non-doxastic) facts, which are true or false in a given world. The valuation tells us which facts hold at which worlds. Finally, the plausibility relation $\leq$ captures the agent's

[5] Cf. similar framework defined in [Ben04].

[6] Having a plausibility frame $(S, \leq)$ Baltag and Smets call them S-propositions.

(conditional) beliefs about the world; if e.g., the agent was given the information that the world is either w or v, she would believe it to be the most plausible of the two. So, if $w < v$, the agent would believe the actual world was w; if $v < w$, she would believe it was v; otherwise (if $w \cong v$), the agent would be indifferent about the two alternatives, she would not be able to decide to believe any one alternative rather than the other.

The existence of minimal elements in any non-empty subset is simply the natural extension of the above setting to general conditional beliefs, not only conditions involving two states. More specifically, for any possible condition $P \subseteq W$ about a system W, the W-proposition $Min_{\leq}P$ is simply a way to encode everything that the agent would believe about the current state of the system W, if she was given the information that the state satisfied condition P.

This brings us to the definition of what it means for a formula $\varphi \in \mathcal{L}_{\Phi}$ to be true at a world w in a single-agent plausibility model M.

Definition 3.8 (Satisfaction relation). Let $M = (W, \leq, \pi)$ be a single-agent plausibility model. The *satisfaction relation* for the propositional connectives is standard (see Definition 1.3), and for the conditional belief is given by, for $w \in W$, $\varphi \in \mathcal{L}_{\Phi}$,

$$M, w \models B^{\psi}\varphi \ \text{ iff } \ M, v \models \varphi \text{ for all } v \in \mathrm{Min}_{\leq}[\![\psi]\!],$$

where $[\![\psi]\!] := \{w \in W : M, w \models \psi\}$. We write $M \models \varphi$ to mean $M, w \models \varphi$ for all $w \in W$. Further, $\models \varphi$ (φ is valid) means that $M \models \varphi$ for all models M.

3.2.1 Saps: Update Mechanism

A belief update is a dynamic form of belief revision, meant to capture the actual change of beliefs induced by learning; the updated belief is about the state of the world as it is after the update. While in the case of (conditional) beliefs the models were kept unchanged, now we have to allow for belief updates that change the original model.

With the setting given by Kripke models, the idea for 'dynamic beliefs' is to use the same type of formalism that was used to model 'static beliefs'. Thus, we model (epistemic or doxastic) actions in essentially the same way as (epistemic) states. In the context of our richer doxastic-plausibility structures, we introduce plausibility preorders on actions and develop a notion of action plausibility models.

As we shall see further (for ranking functions and possibility measures), this approach is usually based on a quantitative interpretation. However, Baltag and Smets' notion of update product is a purely qualitative one. Their anti-lexicographic ordering of pairs consisting of input-state and action, gives priority to the new, incoming information (i.e., to actions). This choice is justified by interpreting the action plausibility model as representing the agent's 'incoming' belief, i.e., the belief-updating event which 'performs' the update by 'acting' on the prior beliefs.

Definition 3.9. An *action plausibility model* is a tuple $(\Sigma, \leq, Pre)$ consisting of a single-agent plausibility space $(\Sigma, \leq)$ together with a *precondition map* $Pre : \Sigma \to \mathrm{Prop}$ associating to each element of Σ some doxastic proposition $Pre(\sigma)$. We call the elements of Σ *(basic) doxastic actions*, preordering $\leq$ the *action plausibility relation*, and $Pre(\sigma)$ the *precondition* of action σ.

The basic actions $\sigma \in \Sigma$ are taken to represent deterministic belief-revising actions of a particularly simple nature. Intuitively, the precondition defines the domain of applicability of action σ; it can be executed on a state w if and only if w satisfies its precondition. The relation $\leq$ gives the agent beliefs about which actions are more plausible than others. "But this should be interpreted as beliefs about changes, that encode changes of beliefs." ([BaS08, p. 41]) In this sense, we use such 'beliefs about actions' as a way to represent doxastic changes; the information about how the agent changes her beliefs is captured by the action plausibility relation. So we read $\sigma < \sigma'$ as saying that if an agent is informed that either σ or σ' is currently happening, then she cannot distinguish between the two, but she believes that σ is in fact happening. Notice that we only deal here with pure belief changes, i.e., actions that do not change the 'ontic' facts of the world, but only the agent's beliefs.

Now we have both – static and dynamic – components and we can proceed to the action-priority update.

Definition 3.10. Let $M = (W, \leq, \pi)$ be a single-agent plausibility model and let $\mathbf{\Sigma} = (\Sigma, \leq, Pre)$ be an action plausibility model. We define their *update product* to be the single-agent plausibility model $M \otimes \mathbf{\Sigma} = (W \otimes \Sigma, \leq, \pi')$, where

1. $W \otimes \Sigma = \{(w, \sigma) \in W \times \Sigma : M, w \models Pre(\sigma)\}$.
2. $\pi'(p) = \{(w, \sigma) \in W \times \Sigma : w \in \pi(p)\}$.
3. $(w, \sigma) \leq (w', \sigma')$ iff either $\sigma < \sigma'$, or else $\sigma \cong \sigma'$ and $w \leq w'$.[7]

Basic actions in this action model are assumed to be deterministic, that is, for a given input-state and a given action, there can only be at most one output-state. More specifically, we select the pairs which are consistent, in the sense that the input-state satisfies the precondition of the action. The updated valuation is essentially given by the original valuation from the input-state model, which expresses the fact that we only consider there 'purely doxastic' actions, i.e., pure belief changes which do not affect the 'ontic' facts of the world (captured here by atomic sentences). The updated plausibility relation is indeed anti-lexicographic preorder relation induced on pairs $(w, \sigma) \in W \times \Sigma$ by the preorders on M and on $\mathbf{\Sigma}$. In other words, the updated plausibility order gives priority to the action plausibility relation (thus resulting in the action-priority update), and apart from this it keeps the original order on states.

In the special case of (truthful) public announcement $!P$, there is only one action in the action model $\mathbf{\Sigma}$ with precondition P. The resulting saps $M \otimes \mathbf{\Sigma}$ consists only of the p-worlds, since they satisfy the proposition P, and the new ordering is the original one restricted to the new domain. Notice that in this case the agent actually gains some knowledge, not just belief (i.e., the update mechanism processes public announcement as designed).

3.2.2 Saps: Equivalence with cps

As the following proposition shows by adding certain properties to a conditional plausibility space we can recover a single-agent plausibility space.[8] The obvious obstacle is that a cps is defined on Popper algebra, with its unconditional version being defined on an algebra, whereas a single-agent plausibility space is defined on

[7] We use the same notation for all three orderings, since their roles are rather obvious.

[8] To show that saps's are, indeed, instances of cps we can use the idea for a preference ordering in [FrH01, p. 9].

worlds. Once we make sure that all singletons $\{w\}$ for $w \in W$ are measurable, we also need to make sure that they (resp. their plausibilities) are well-preordered. Since a preorder is already induced by $\leq_D$ on $Pl(w)$ for $w \in W$, we only need to add a requirement that every non-empty subset consisting of these plausibilities $Pl(w)$ has maximal elements.[9] After that we add some extra properties in order for the cps to deal with arbitrary (plausible) public announcement and to process implausibility the same way as a saps does.

Proposition 3.11. *Let $(W, \mathcal{F}, \mathcal{F}', Pl)$ be a conditional plausibility space (defined by CPl1–CPl5; see Definition 2.7). In order to recover a single-agent plausibility space $(W, \leq)$ the following properties need to be added:*

1. $\mathcal{F} = 2^W$.
2. $D_W = \{Pl(V) : V \in \mathcal{F}\}$ *is totally ordered by* $\leq_D$*: either* $Pl(A) \leq_D Pl(B)$ *or* $Pl(B) \leq_D Pl(A)$ *for all sets* $A, B \in \mathcal{F}$.[10]
3. $Pl(A \cup B) = \max(Pl(A), Pl(B))$ *for all sets* $A, B \in \mathcal{F}$.
4. *The given cps* $(W, \mathcal{F}, \mathcal{F}', Pl)$ *is standard.*
5. $Pl(A) = \bot$ *iff* $A = \emptyset$.

Proof. For all $w \in W$ we have a plausibility value $Pl(\{w\} \mid W)$, also written as $Pl(w)$, which is given by $Pl : 2^W \times \mathcal{F}' \to D$, where $W \in \mathcal{F}'$, because $\mathcal{F}'$ is closed under supersets in $\mathcal{F}$. Then we take an order $\leq$ on W to be the inverse of $\leq_D$ on the corresponding elements in D, that is, $v \leq w$ iff $Pl(w) \leq_D Pl(v)$ for all $v, w \in W$. It is easy to see that $\leq$ is a well-preorder on W. It is reflexive and transitive, because its inverse is a partial order on D[11]. Notice that $\{Pl(w) : w \in W\} \subseteq D_W$ which is totally ordered by $\leq_D$. From the first three properties we have that $Pl(A) = \max_{w \in A} Pl(w)$ for all $A \in \mathcal{F}$. Therefore every non-empty subset of W has $\leq$-minimal elements, because the corresponding subset

[9] Notice that the order on worlds and the one on their plausibilities are in a certain sense reversed; while the best worlds in a saps are the minimal ones, in a cps they have the maximal plausibility.

[10] Notice that by CPl5 it is also the case that either $Pl(A \mid U) \leq_D Pl(B \mid U)$ or $Pl(B \mid U) \leq_D Pl(A \mid U)$ for all $U \in \mathcal{F}'$, i.e., Pl_U places a total order on sets of $\mathcal{F}$.

[11] $w \leq w$ iff $Pl(w) \leq_D Pl(w)$, and if $u \leq v$ and $v \leq w$ then also $Pl(v) \leq_D Pl(u)$ and $Pl(w) \leq_D Pl(v)$ and by transitivity $Pl(w) \leq_D Pl(u)$ which is iff $u \leq w$.

of plausibilities in D has $\leq_D$-maximal elements. Thus, the resulting plausibility space $(W, \leq)$ is in fact a single-agent plausibility space.

The fourth property ensures that we can update the cps with every plausible public announcement. It shall become clearer below together with the purpose of the fifth property. □

Our aim is to show that a cps on W with the properties stated above is such that if it is equivalent to a saps on W then these spaces will remain equivalent even after they both undergo public announcement $!P$. To be precise, by equivalent we mean that for a fixed valuation π their corresponding structures continue satisfying the same formulas after the update.

Let the updated versions be $(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P, \pi_P)$ for a cps and $(W_P, \leq_P, \pi_P)$ for a saps, where $W_P = W \cap P$, and π_P and $\leq_P$ are obtained from the original π and $\leq$ by restricting them to the new domains W_P, and $W_P \times W_P$, respectively. We want to make sure that for all $\varphi \in \mathcal{L}_\Phi$: $(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P, \pi_P) \models \varphi$ if and only if $(W_P, \leq_P, \pi_P) \models \varphi$. Propositional cases (i.e., atoms, negation and conjunction) are straightforward since we have the same universe W_P and the same valuation π_P. The only case left to be checked is conditional belief $B^\psi\varphi$. However, a statement such as $B^\psi\varphi$ is determined by the global plausibility measure. Thus, the set of worlds that satisfy $B^\psi\varphi$ is either the empty set or W:

(i) A single-agent plausibility model:
$(W_P, \leq_P, \pi_P) \models B^\psi\varphi$ iff $(W_P, \leq_P, \pi_P), v \models \varphi$ for all $v \in \mathrm{Min}_\leq[\![\psi]\!]$.

(ii) A conditional plausibility structure:
$(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P, \pi_P) \models B^\psi\varphi$ iff $Pl_P([\![\psi]\!]) = \bot$ or $Pl_P([\![\varphi \wedge \psi]\!]) >_D Pl_P([\![\neg\varphi \wedge \psi]\!])$.

Therefore, in order to make sure that both cps and saps on W (resp. their structures) satisfy the same formulas after public announcement, we only need to make sure that they satisfy the same conditional beliefs. For that reason let us first define an ordering on sets of worlds for a saps.

Definition 3.12. Let $(W, \leq)$ be a single-agent plausibility space and $V, U \subseteq W$. We define that $V \leq^S U$ iff either there exist $v \in Min_\leq V$ and $u \in Min_\leq U$ such that $v \leq u$, or if $U = \emptyset$. We write $V <^S U$ iff $V \leq^S U$ and $U \not\leq^S V$, and $V \cong^S U$ iff both $V \leq^S U$ and $U \leq^S V$.

Clearly, $\leq^S$ is a well-preorder on sets, and it extends $\leq$: $v \leq u$ iff $\{v\} \leq^S \{u\}$. It is also the case that $<^S$ extends $<$.

Now we can redefine the satisfaction condition for conditional beliefs on a saps from above as: $(W_P, \leq_P, \pi_P) \models B^{\psi}\varphi$ iff $[\![\psi]\!] = \emptyset$ or $[\![\varphi \wedge \psi]\!] <^S [\![\neg\varphi \wedge \psi]\!]$. The following proposition proves that it is well defined.

Proposition 3.13. *Given a single-agent plausibility model* $(W, \leq, \pi)$ *and formulas* φ, $\psi \in \mathcal{L}_\Phi$, *we claim that the following are equivalent:*

(i) $(W, \leq, \pi), v \models \varphi$ *for all* $v \in Min_{\leq}[\![\psi]\!]$.

(ii) $[\![\psi]\!] = \emptyset$ *or* $[\![\varphi \wedge \psi]\!] <^S [\![\neg\varphi \wedge \psi]\!]$.

Proof. (i)→(ii): We want to prove that either $[\![\psi]\!] = \emptyset$ or $[\![\varphi \wedge \psi]\!] <^S [\![\neg\varphi \wedge \psi]\!]$, which by Definition 3.12 is iff either there exist $v \in Min_{\leq}[\![\varphi \wedge \psi]\!]$ and $u \in Min_{\leq}[\![\neg\varphi \wedge \psi]\!]$ such that $v < u$, or $[\![\neg\varphi \wedge \psi]\!] = \emptyset$ (and $[\![\varphi \wedge \psi]\!] \neq \emptyset$). However, from (i) we have that the minimal $[\![\psi]\!]$-worlds satisfy also φ. If $[\![\psi]\!] = \emptyset$ then we are done. Otherwise, let w be one of these worlds. Notice that w is exactly the witness we need, since $w \in Min_{\leq}[\![\varphi \wedge \psi]\!]$ and either $[\![\neg\varphi \wedge \psi]\!] = \emptyset$, or there is some $u \in Min_{\leq}[\![\neg\varphi \wedge \psi]\!]$, in which case it must be that $w < u$, because all the minimal $[\![\psi]\!]$-worlds are in $[\![\varphi \wedge \psi]\!]$. Thus, $[\![\varphi \wedge \psi]\!] <^S [\![\neg\varphi \wedge \psi]\!]$.
(ii)→(i): We want to prove that $(W, \leq, \pi), w \models \varphi$ for all $w \in Min_{\leq}[\![\psi]\!]$. However, from (ii) we have that either $[\![\psi]\!] = \emptyset$, or $[\![\varphi \wedge \psi]\!] <^S [\![\neg\varphi \wedge \psi]\!]$. The first case is trivial. In the second case either there exist $v \in Min_{\leq}[\![\varphi \wedge \psi]\!]$ and $u \in Min_{\leq}[\![\neg\varphi \wedge \psi]\!]$ such that $v < u$, or $[\![\neg\varphi \wedge \psi]\!] = \emptyset$ (and $[\![\varphi \wedge \psi]\!] \neq \emptyset$). In the former case the minimal $[\![\psi]\!]$-worlds must be all in $[\![\varphi \wedge \psi]\!]$ and as such they satisfy φ. In the latter case there are only $[\![\psi]\!]$-worlds which satisfy φ. □

Let us define an analogue to default-equivalence between two plausibility spaces.

Definition 3.14. We say that a conditional plausibility space (W, Pl) (resp. conditional plausibility structure (W, Pl, π)) and a single-agent plausibility space $(W, \leq)$ (resp. single-agent plausibility model $(W, \leq, \pi)$) are *default-equivalent* if for all disjoint subsets A and B of W it holds $Pl(A) <_D Pl(B)$ if and only if $B <^S A$.

Clearly, if a structure (W, Pl, π) and a model $(W, \leq, \pi)$ are default-equivalent, then they satisfy the same conditional beliefs.

Proposition 3.15. *Any conditional plausibility space* $(W, \mathcal{F}, \mathcal{F}', Pl)$ *satisfying the conditions in Proposition 3.11 is default-equivalent to the recovered single-agent plausibility space* $(W, \leq)$.

Proof. We want to show that for all disjoint subsets A and B of W, we have $Pl(A) <_D Pl(B)$ in the cps if and only if $B <^S A$ in the saps. It is rather obvious that our construction of the saps in Proposition 3.11 leads to this conclusion.

The first property of the proposition makes sure that all subsets of W are measurable and the second property subsequently enables their comparison. The third property regulates $Pl(A)$ so that the correlation on worlds can be lifted onto sets of worlds. Recall that we have constructed the saps via the relation $v \leq w$ iff $Pl(w) \leq_D Pl(v)$ on the worlds $w \in W$. Therefore, the minimal elements in the saps have maximal plausibility in the cps, and vice versa. Notice that subsets $A \subseteq W$ in the saps are ordered according to their minimal elements in $Min_{\leq} A$, and their plausibilities in the cps are ordered according to the elements with maximal plausibility, since $Pl(A) = \max_{w \in A} Pl(w)$. What is more, the fifth property ensures that $Pl(\emptyset) <_D Pl(B)$ iff $B <^S \emptyset$. Thus, it follows that the cps and the recovered saps are default-equivalent. □

Theorem 3.16. *Any conditional plausibility space* $(W, \mathcal{F}, \mathcal{F}', Pl)$ *satisfying the conditions in Proposition 3.11 is such that the following property holds:*
If $Pl(P) \neq \bot$ *and the cps is default-equivalent to a saps* $(W, \leq)$, *then the new cps* $(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P)$ *is default-equivalent to the saps obtained by updating* $(W, \leq)$ *with public announcement* $!P$. *Moreover, the new cps also satisfies the conditions in Proposition 3.11.*

Proof. First of all, note that according to Proposition 3.11 (fourth property) the cps $(W, \mathcal{F}, \mathcal{F}', Pl)$ is standard. Hence, $W \cap P \in \mathcal{F}'$ and we can apply the function R_{PA} of public announcement in Definition 3.2. The single-agent plausibility space (its model, to be precise), on the other hand, is modified by the update mechanism as described in Subsection 3.2.1.

Further, we need to make sure that the updated cps $(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P)$ also satisfies the conditions in Proposition 3.11:

1. $\mathcal{F}_P = \{U \cap P : U \in \mathcal{F}\}$, and so if $\mathcal{F} = 2^W$, then $\mathcal{F}_P = 2^{W_P}$.

2. $D_P = \{Pl_P(V) : V \in \mathcal{F}\}$ is indeed totally ordered by $\leq_D$. We have that either $Pl(A) \leq_D Pl(B)$ or $Pl(B) \leq_D Pl(A)$ for all sets $A, B \in \mathcal{F}$. Since $\mathcal{F}$ is an algebra, $A \cap P, B \cap P \in \mathcal{F}_P \subseteq \mathcal{F}$. Thus, notice that by CPl5 (see Definition 2.7) it is also the case that either $Pl(A \mid P) \leq_D Pl(B \mid P)$ or $Pl(B \mid P) \leq_D Pl(A \mid P)$, i.e., Pl_P also places a total order on sets of $\mathcal{F}$.

3. Since $A \cap P, B \cap P \in \mathcal{F}_P \subseteq \mathcal{F}$ we have $Pl((A \cap P) \cup (B \cap P)) = \max(Pl(A \cap P), Pl(B \cap P))$. Thus, it must be either $Pl((A \cup B) \cap P) = Pl(A \cap P)$ which by CPl5 is iff $Pl((A \cup B) \mid P) = Pl(A \mid P)$, or $Pl((A \cup B) \cap P) = Pl(B \cap P)$ which is iff $Pl((A \cup B) \mid P) = Pl(B \mid P)$. Indeed it depends on a relation between $Pl(A \cap P)$ and $Pl(B \cap P)$. Notice that by CPl5 also $Pl(A \cap P) \leq Pl(B \cap P)$ iff $Pl(A \mid P) \leq Pl(B \mid P)$. It follows that $Pl((A \cup B) \mid P) = \max(Pl(A \mid P), Pl(B \mid P))$, which can be written as $Pl_P(A \cup B) = \max(Pl_P(A), Pl_P(B))$.

4. If $\mathcal{F}' = \{U : Pl(U \mid W) \neq \bot\}$ then $\mathcal{F}'_P = \{U \cap P : Pl(U \mid W) \neq \bot$ & $Pl(U \mid P) \neq \bot\}$, which (by CPl4) can be also written as $\mathcal{F}'_P = \{U \cap P : Pl(U \cap P \mid W \cap P) \neq \bot\}$. Thus, the updated cps $(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P)$ is standard.

5. If $Pl_P(A \cap P) = \bot$ then by CPl5 we have $Pl(A \cap P) = \bot$ and from the properties of the original cps it follows that $A \cap P = \emptyset$. The opposite direction is straightforward.

The next step is to show that if the original plausibility spaces were default-equivalent, then their updated versions are also default-equivalent, i.e., that for all disjoint subsets A and B of W_P we have $Pl_P(A) <_D Pl_P(B)$ iff $B <^S A$. Then, indeed, their corresponding structures satisfy the same conditional beliefs:

(i) A single-agent plausibility model:
$(W_P, \leq_P, \pi_P) \models B^\psi \varphi$ iff $[\![\psi]\!] = \emptyset$ or $[\![\varphi \wedge \psi]\!] <^S [\![\neg\varphi \wedge \psi]\!]$.

(ii) A conditional plausibility structure:
$(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P, \pi_P) \models B^\psi \varphi$ iff $Pl_P([\![\psi]\!]) = \bot$ or $Pl_P([\![\varphi \wedge \psi]\!]) >_D Pl_P([\![\neg\varphi \wedge \psi]\!])$.[12]

[12]Notice that in this case π_P as defined for a saps in Subsection 3.2.1 is in accordance with a valuation π relativised to the new domain W_P.

Recall that for both spaces the ordering on (plausibilities of) sets of worlds is determined by the ordering on (plausibilities of) worlds. We also know that public announcement $!P$ simply eliminates the $\neg p$-worlds, but otherwise leaves the original ordering on the p-worlds. Hence, it is easy to see that both updates with public announcement preserve default-equivalence between the spaces, in particular, between their structures (see the summarising picture below).

It is rather straightforward that these updated structures satisfy the same formulas in $\mathcal{L}_\Phi$, since the propositional cases are obvious and we have also proved the case of conditional belief.

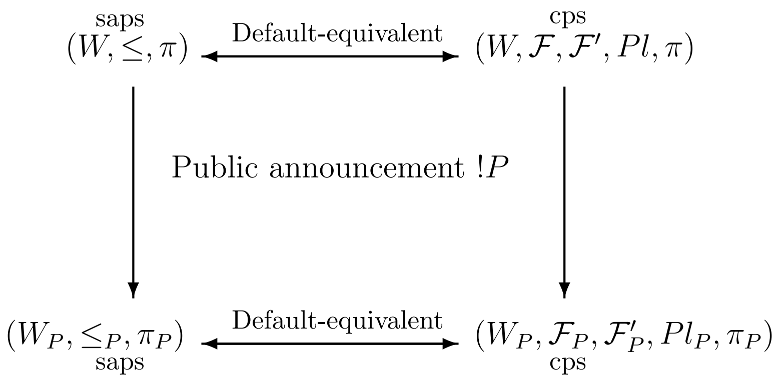

□

3.3 Kappa-ranking

In this section we present similar results for ranking functions as we did in the previous section for single-agent plausibility spaces. First, we introduce the update mechanism with its static and dynamic components, and then the relation between a space for ranking function and a conditional plausibility space after public announcement on both of them.

3.3.1 Kappa-ranking: Update Mechanism

The core idea of the update mechanism for (conditional) ranking functions (also called ordinal (conditional) functions) originally comes from W. Spohn in [Spo88], but the update mechanism introduced below has been defined in [Auc03].[13]

[13] In [Auc07] an analogue of such update mechanism for probability can be found.

With the richer doxastic-plausibility structures, epistemic changes require as their input not only propositions but also their degree of plausibility[14]. We shall note this input as a pair (A, α) where A is the proposition and α the degree of plausibility with which the formula should be believed by an agent after her revision of beliefs.

Definition 3.17. Let κ be a conditional ranking function on W and $A \subseteq W$. The A-part of κ is the function $\kappa(. \mid A)$ defined on A for which for all $w \in A$, $\kappa(w \mid A) = \kappa(w) - \kappa(A)$.

We could say that the A-part of κ is the restriction of κ to A shifted to 0, that is, in such a way that $\kappa(A \mid A) = 0$.

Using this concept, we can now define the conditional ranking function $\kappa_{A,\alpha}$ representing the new state of belief:

Definition 3.18. Let A be a proposition such that $A \neq \emptyset$, $A \neq W$, and α an ordinal[15]. Then we define $\kappa_{A,\alpha}$ the (A, α)-*conditionalisation of* κ as follows:

$$\kappa_{A,\alpha}(w) = \begin{cases} \kappa(w \mid A) & \text{if } w \in A \\ \alpha + \kappa(w \mid \overline{A}) & \text{if } w \in \overline{A} \end{cases}$$

Thus the (A, α)-conditionalisation of κ is the union of the A-part of κ and the $\overline{A}$-part of κ shifted up by α degrees of plausibility. It follows from the definition that $\kappa_{A,\alpha}(A) = 0$ and $\kappa_{A,\alpha}(\overline{A}) = \alpha$. Hence we say that A is believed in $\kappa_{A,\alpha}$ with firmness α. Importantly, notice that the (A, α)-conditionalisation of κ leaves the A-part as well as the $\overline{A}$-part of κ unchanged, they are only shifted relatively to each other.

In order to define the update mechanism in [Auc03], first we need to briefly introduce its static and dynamic components. Since we are only interested in a single-agent case, we provide a simplified version by removing the epistemic relations $\{\sim_j : j \in G\}$, for a finite set of agents G. We also omit the actual world w_0, and the actual action σ_0, respectively, from the corresponding 'pointed' models.

For the static part we define a belief model rather naturally.

[14]Spohn refers to it as a degree of firmness.
[15]For our purposes $\alpha \in \mathbb{N}$ suffices.

Definition 3.19. A *belief model* $M = (W, \kappa, \pi)$ is a tuple where:[16]

1. W is a set of possible worlds (or states) of the model.
2. κ is an operator, ranging on natural numbers from 0 to Max, defined on all the worlds.
3. π is a valuation.

The ranking function κ ranges from 0 to Max, where Max is an arbitrary fixed natural number. The more a world is plausible for the agent, the closer its plausibility value is to 0, and the less plausible a world is the closer it is to Max. The intuition behind introducing a natural number Max is that beyond a certain degree of plausibility the agent cannot distinguish two different worlds of different plausibility.[17]

As for the dynamic part, we define a belief action model following the same pattern as above.

Definition 3.20. A *belief action model* $\mathbf{\Sigma} = (\Sigma, \kappa^*, Pre)$ is a tuple where:[18]

1. Σ is a set of simple actions.
2. κ^* is a function from the set of simple actions to the set of natural numbers ranging from 0 to Max, where at least one of the actions is assigned the plausibility 0.
3. Pre is a function from the set of simple actions to the formulas of $\mathcal{L}_\Phi$ (resp. doxastic propositions).

By simple actions we mean that they cannot be decomposed into 'smaller (sub)actions' whose succession would form the original action. The ranking function κ^* expresses the plausibility preference that the agent has among actions that she cannot objectively distinguish. Pre assigns to each simple action a precondition that a world must fulfill in order for this action to be performed in this world (e.g., One can hardly read a book in the world where there are no books.).

This brings us to the definition of what it means for a formula $\varphi \in \mathcal{L}_\Phi$ to be true at a world w in a belief model M (resp. in a ranking structure).

[16] Originally it is a belief epistemic model $M = (W, \{\sim_j : j \in G\}, \{\kappa_j : j \in G\}, \pi, w_0)$.

[17] However, the main reasons for introducing Max were, indeed, technical.

[18] Originally it is a belief epistemic action model $\mathbf{\Sigma} = (\Sigma, \sim_j, \kappa^*_j, Pre, \sigma_0)$.

Definition 3.21 (Satisfaction relation)**.** Let $M = (W, \kappa, \pi)$ be a belief model (resp. a ranking structure). The *satisfaction relation* for the propositional connectives is standard (see Definition 1.3), and for the conditional belief is given by, for $w \in W$, $\varphi \in \mathcal{L}_\Phi$,

$$M, w \models B^{\psi}\varphi \;\text{ iff }\; \kappa([\![\psi]\!]) = \infty \text{ or } \kappa([\![\varphi \wedge \psi]\!]) < \kappa([\![\neg\varphi \wedge \psi]\!]),$$

where $[\![\psi]\!] := \{w \in W : M, w \models \psi\}$.

We have now defined two main components. First, the belief model M as a formal counterpart of the way an actual situation s is perceived by the agent according to her beliefs and knowledge. Second, the belief action model $\mathbf{\Sigma}$ as a formal counterpart of the way an actual action a is perceived by the agent according to her beliefs and knowledge. In reality the agent updates her knowledge and beliefs according to these two pieces of information: situation s and action a, giving rise to a new actual situation $s \times a$. Formally, this update is determined by a 'mathematical model' $\otimes$ such that $M \otimes \mathbf{\Sigma}$ is a new belief model and a formal counterpart of $s \times a$.

Definition 3.22. Given a belief model $M = (W, \kappa, \pi)$ and a belief action model $\mathbf{\Sigma} = (\Sigma, \kappa^*, Pre)$ we define their *update product* to be the belief model $M \otimes \mathbf{\Sigma} = (W \otimes \Sigma, \kappa', \pi')$ where:

1. $W \otimes \Sigma = \{(w, \sigma) \in W \times \Sigma : M, w \models Pre(\sigma)\}$;

2. $\pi'(p) = \{(w, \sigma) \in W \times \Sigma : w \in \pi(p)\}$;

3. $\kappa'(w, \sigma) = Cut_M(\kappa^*(\sigma) + \kappa(w) - \kappa(Pre(\sigma)))$ where

   $\kappa(Pre(\sigma)) = \min\{\kappa(v) : M, v \models Pre(\sigma)\}$ and

   $$Cut_M(x) = \begin{cases} x & \text{if } 0 \le x \le Max \\ Max & \text{if } x > Max \end{cases}$$

A new possible world (w, σ) in the resulting model corresponds to the resulting situation of performing the action corresponding to σ in the world corresponding to w. However, the action σ can be performed in w only if the precondition $Pre(\sigma)$ of the action σ is satisfied in the world w. Actions cannot be performed in an arbitrary world, they presuppose some 'material' preconditions in the world. We essentially take the same valuation as the one of the input model. This means that here we consider only the actions which do not change the facts (i.e., 'purely doxastic' actions).

Recall that a domain and a valuation in a single-agent plausibility model are updated the exact same way. However, that is not the case for an order $\leq$ and a ranking function κ, respectively.

The core of the update is $\kappa^*(\sigma)+\kappa(w)-\kappa(Pre(\sigma))$. Cut_M is only a technical device assuring that the new plausibility assignment fits in the range of the plausibility scale of the new belief model. $\kappa(w)$ is plausibility for the agent that w is the actual world and $\kappa^*(\sigma)$ is plausibility for the agent that σ is the actual action taking place in w.

Let us take a better look at the part of this 'conditioning': $-\kappa(Pre(\sigma)) = -\min\{\kappa(v) : M, v \models Pre(\sigma)\}$. The idea is to relativise the ordinal assignment to the relevant worlds, that is, the worlds where σ can take place. Indeed, the former ordinal assignment κ made sense only if we considered all the worlds that may correspond for the agent to the actual situation w. However, since the action σ is taking place, the agent needs to restrict her attention only to the worlds where σ can take place, and obviously the other ones do not play a role any more. So the agent must relativise her plausibility ordering to this set by rescaling the ordinal assignment in order to start again and deal with the action σ with a self defined plausibility ordering.

As for a single-agent plausibility space, in the special case of (truthful) public announcement $!P$ we obviously have only one action in the belief action model $\mathbf{\Sigma}$ with precondition P and plausibility $\kappa^*(\sigma) = 0$. It means that in the updated belief model $M \otimes \mathbf{\Sigma}$ we have reduced worlds of M to the p-worlds, since $W \otimes \Sigma = \{(w, \sigma) \in W \times \{\sigma\} : w \in P\}$. Moreover, the update results in a new ranking which has shifted the p-worlds to 0 (cf. Definition 3.17): $\kappa'(w, \sigma) = \kappa(w) - \min\{\kappa(v) : v \in P\} = \kappa(w) - \kappa(P)$.[19]

[19] If we considered a conditional ranking function κ, in the case of public announcement $!P$ we could take $\kappa'(w, \sigma) = \kappa(w \mid P)$.

3.3.2 Kappa-ranking: Equivalence with cps

As we have already seen in the previous chapters, ranking functions are indeed instances of a (conditional) plausibility measure. However, what properties do we need to add to a cps to actually recover a κ-ranking[20]? Recall that a κ-ranking is defined on 2^W, ranges over $\mathbb{N}^* = \mathbb{N} \cup \infty$ with 0 being the best and ∞ denoting the worst plausibility (it is usually explained in terms of surprise), and plausibility of a union is determined by the plausibilities of its components (see Definition 1.8). Notice that the following proposition lists the same properties as Proposition 3.11 except for the last one dealing with implausibility of non-empty sets. We discuss the matter at the end of this chapter.

Proposition 3.23. *Let $(W, \mathcal{F}, \mathcal{F}', Pl)$ be a conditional plausibility space (defined by CPl1–CPl5; see Definition 2.7). In order to recover a ranking function κ on W the following properties need to be added:*

1. $\mathcal{F} = 2^W$.
2. $D_W = \{Pl(V) : V \in \mathcal{F}\}$ *is totally ordered by* $\leq_D$*: either* $Pl(A) \leq_D Pl(B)$ *or* $Pl(B) \leq_D Pl(A)$ *for all sets* $A, B \in \mathcal{F}$.
3. $Pl(A \cup B) = \max(Pl(A), Pl(B))$ *for all sets* $A, B \in \mathcal{F}$.
4. *The given cps* $(W, \mathcal{F}, \mathcal{F}', Pl)$ *is standard.*

Proof. Let κ be a function mapping subsets of W to $\mathbb{N}^* = \mathbb{N} \cup \infty$ such that for all $A, B \subseteq W$ we define $\kappa(A) \leq \kappa(B)$ iff $Pl(B) \leq_D Pl(A)$. According to the first two properties this way κ is well-defined on all subsets of W as desired. The third property determines plausibility of a union as the maximal plausibility of the two components. Clearly, it follows that for plausibility κ we have $\kappa(A \cup B) = \min(\kappa(A), \kappa(B))$. Last but not least, for all $A \subseteq W$ we require that $\kappa(A) = 0$ iff $Pl(A) = \top$ and $\kappa(A) = \infty$ iff $Pl(A) = \bot$. Notice that it must be the case that $\min_{w \in W} \kappa(w) = 0$. Hence, the resulting plausibility measure κ is in fact a ranking function.

The fourth property ensures that we can update the cps with every plausible public announcement. It shall become clearer below. □

[20] We use the notion of κ-ranking to refer to the ranking function κ as well as the corresponding space $(W, 2^W, \kappa)$.

As in the case of a single-agent plausibility space, we want to show that if a conditional plausibility structure on W with the properties stated above and a ranking structure on W satisfy the same formulas they continue to do so even after public announcement. For the same reason as before (see Subsection 3.2.2), we only need to focus on conditional beliefs.

Let us define an analogue to default-equivalence between two plausibility spaces.

Definition 3.24. We say that a conditional plausibility space (W, Pl) (resp. conditional plausibility structure (W, Pl, π)) and a κ-ranking (W, κ) (resp. a ranking structure[21] (W, κ, π)) are *default-equivalent* if for all disjoint subsets A and B of W it holds $Pl(A) <_D Pl(B)$ if and only if $\kappa(B) < \kappa(A)$.

Clearly, if the structures (W, Pl, π) and (W, κ, π) are default-equivalent, then they satisfy the same conditional beliefs.

Proposition 3.25. *Any conditional plausibility space $(W, \mathcal{F}, \mathcal{F}', Pl)$ satisfying the conditions in Proposition 3.23 is default-equivalent to the recovered κ-ranking $(W, 2^W, \kappa)$.*

Proof. We want to show that for all disjoint subsets A and B of W, we have $Pl(A) <_D Pl(B)$ if and only if $\kappa(B) < \kappa(A)$. However, their default-equivalence follows directly from the construction of κ-ranking in Proposition 3.23. □

Now we can proceed to the main theorem. The idea is that if a cps is default-equivalent to a κ-ranking and they are updated with plausible public announcement $!P$, the updated versions of these spaces are also default-equivalent. Whereas the cps is updated according to Definition 3.2, the κ-ranking is modified by the update mechanism defined in Subsection 3.3.1 and results in a new space $(W_P, 2^{W_P}, \kappa_P)$. What is more, the properties in Proposition 3.23 are also preserved.

Theorem 3.26. *Any conditional plausibility space $(W, \mathcal{F}, \mathcal{F}', Pl)$ satisfying the conditions in Proposition 3.23 is such that the following property holds:*
If $Pl(P) \neq \perp$ and the cps is default-equivalent to a κ-ranking $(W, 2^W, \kappa)$, then the new cps $(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P)$ is default-equivalent to the κ-ranking obtained by

[21] In the previous subsection we called it a belief model.

updating $(W, 2^W, \kappa)$ *with public announcement* $!P$*. Moreover, the new cps also satisfies the conditions in Proposition 3.23.*

Proof. As expected the proof follows almost identical path to the one in the case of a single-agent plausibility space in Theorem 3.16. Note that since the cps is standard we can apply the function R_{PA} of public announcement in Definition 3.2 and get the new cps $(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P)$.

On the other hand, the ranking function κ is modified by the update mechanism as described in Subsection 3.3.1: $\kappa_P(w) = \kappa(w) - \kappa(P)$. This time, however, we omitted the limitation Max and included the value ∞ instead. The mechanism works as it did before with only one exception – if a world has plausibility of value ∞, it is also assigned plausibility ∞ as the resulting value (if $\kappa(w) = \infty$ then $\kappa_P(w) = \infty$), even in the case when all the finite values are being shifted towards 0.[22] Notice also that it cannot be the case that $\kappa(P) = \infty$, since the original spaces are default-equivalent and it is given that $Pl(P) \neq \bot$.

Further borrowing the arguments from Theorem 3.16, we can conclude that the new cps satisfies the conditions in Proposition 3.23, since they have already been proved. The only thing left to show is that both spaces remain default-equivalent after the update with public announcement $!P$, i.e., for all disjoint subsets A and B of W_P we have $Pl_P(A) <_D Pl_P(B)$ iff $\kappa_P(B) < \kappa_P(A)$. Then, indeed, their corresponding structures satisfy the same conditional beliefs:

(i) A ranking structure:
$(W_P, 2^{W_P}, \kappa_P, \pi_P) \models B^\psi\varphi$ iff $\kappa_P([\![\psi]\!]) = \infty$ or $\kappa_P([\![\varphi\wedge\psi]\!]) < \kappa_P([\![\neg\varphi\wedge\psi]\!])$.

(ii) A conditional plausibility structure:
$(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P, \pi_P) \models B^\psi\varphi$ iff $Pl_P([\![\psi]\!]) = \bot$ or $Pl_P([\![\varphi \wedge \psi]\!]) >_D Pl_P([\![\neg\varphi \wedge \psi]\!])$.[23]

Recall that for both spaces the ordering on plausibilities of sets of worlds is determined by the ordering on plausibilities of worlds. Public announcement $!P$ eliminates the $\neg p$-worlds and what is more, in the case of a κ-ranking it also shifts all the 'surviving' p-worlds towards 0 (except for those with plausibility ∞). However, the original orderings on the p-worlds remain the same in both

[22] In the spirit of the motto: If something is '∞-implausible', it cannot be recovered.

[23] Notice that in this case π_P as defined for a κ-ranking in Subsection 3.3.1 is in accordance with a valuation π relativised to the new domain W_P.

spaces. And as such both updates with public announcement preserve default-equivalence between the spaces, in particular, between their structures (see the summarising picture below).

It is rather straightforward that these updated structures satisfy the same formulas in $\mathcal{L}_\Phi$, since the propositional cases are obvious and we have just proved the case of conditional belief.

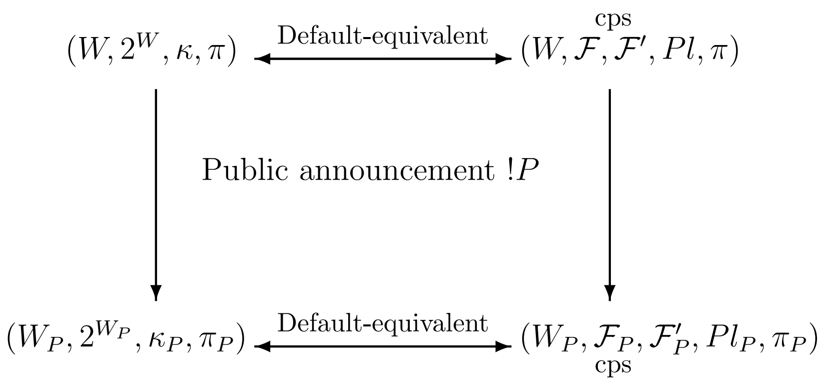

□

3.4 Possibility

In this section we present similar results for possibility measures as we did in the two previous sections for single-agent plausibility spaces, and ranking functions, respectively.

3.4.1 Possibility: Update Mechanism

Let us slightly modify conditioning in possibility and briefly introduce its potential. There are several conditioning methods in possibility theory, and here we adopt the one in [DuP93] (as suggested in [MaL08]).

We take $Poss$ to be a conditional possibility measure satisfying CPoss1–CPoss4 as stated in Definition 2.6, but in CPoss4 we substitute $\times$ for min. It results in the following form of conditioning:

$$Poss(U_1 \mid U_2 \cap U_3) = \frac{Poss(U_1 \cap U_2 \mid U_3)}{Poss(U_2 \mid U_3)}, \tag{3.1}$$

with an additional condition that $Poss(U_1 \mid U_2 \cap U_3) = 0$ if $Poss(U_2 \mid U_3) = 0$.

A counterpart of Spohn's (A, α)-conditionalisation was suggested in possibility theory in [DuP93] (resp. in [MaL08]) for change of an agent's current belief $Poss$ when the new evidence claims that $Poss'(A) = 1$ and $Poss'(\overline{A}) = 1 - \alpha$ (which implies that $Nec'(A) = \alpha$).[24]

Definition 3.27. Let A be a proposition such that $A \neq \emptyset$, $A \neq W$, and $\alpha \in [0, 1]$. Then we define $Poss_{A,\alpha}$ the (A, α)*-conditionalisation of* $Poss$ as follows:

$$Poss_{A,\alpha}(w) = \begin{cases} Poss(w \mid A) & \text{for } w \in A \\ (1-\alpha)Poss(w \mid \overline{A}) & \text{for } w \in \overline{A} \end{cases} \tag{3.2}$$

We can derive that $Poss(w \mid A) = Poss(w)/Poss(A)$ from (3.1) with U_1 being a singleton $\{w\}$, $U_2 = A$, $U_3 = W$.

In order to state a theorem similar to the previous ones (Theorem 3.16 and Theorem 3.26, respectively) we are missing an actual update mechanism for possibility measures. However, for our purposes we can abstract the core idea from the counterpart of Spohn's (A, α)-conditionalisation introduced above, the same way it has been done in the update mechanism for ranking functions (cf. Subsection 3.3.1).

In the case of (possible) public announcement $!P$ we are left only with the upper branch in (3.2) simply conditioning the possibility of a given p-world to the possibility of the set P of all the p-worlds: $Poss(w \mid P) = Poss(w)/Poss(P)$. Obviously, this way possibility of the individual p-worlds can only increase (except for those of possibility 0 which remain the same), but the original ordering on them is preserved.

This brings us to the definition of what it means for a formula $\varphi \in \mathcal{L}_\Phi$ to be true at a world w in a possibility structure PS.

Definition 3.28 (Satisfaction relation)**.** Let $PS = (W, Poss, \pi)$ be a possibility structure. The *satisfaction relation* for the propositional connectives is standard (see Definition 1.3), and for the conditional belief is given by, for $w \in W$, $\varphi \in \mathcal{L}_\Phi$,

$$PS, w \models B^\psi \varphi \;\text{ iff }\; Poss([\![\psi]\!]) = 0 \text{ or } Poss([\![\varphi \wedge \psi]\!]) > Poss([\![\neg\varphi \wedge \psi]\!]),$$

where $[\![\psi]\!] := \{w \in W : PS, w \models \psi\}$.

[24] Recall that $Poss(A)$ estimates the degree an agent believes the true world can be in A while $Nec(A) = 1 - Poss(\overline{A})$ estimates the degree the agent believes the true world should be necessarily in A.

3.4.2 Possibility: Equivalence with cps

As we have already seen in the previous chapters, possibility measures are indeed instances of a (conditional) plausibility measure. However, what properties do we need to add to a cps to actually recover a possibility measure? Recall that $Poss$ is defined on 2^W, ranges over $[0, 1]$ with 0 being the worst and 1 being the best possibility, and possibility of a union is determined by the possibilities of its components (see Definition 1.9). Notice that we select the same properties as for ranking functions in Proposition 3.23, which in return are those for saps's in Proposition 3.11 except for the last one dealing with implausibility of non-empty sets. We discuss the similarities at the end of this chapter.

Proposition 3.29. *Let $(W, \mathcal{F}, \mathcal{F}', Pl)$ be a conditional plausibility space (defined by CPl1–CPl5; see Definition 2.7). In order to recover a possibility measure Poss on W exactly the properties from Proposition 3.23 need to be added.*

Proof. Let $Poss$ be a function mapping subsets of W to $[0, 1]$ such that for all $A, B \subseteq W$ we define $Poss(A) \leq Poss(B)$ iff $Pl(A) \leq_D Pl(B)$. Again, the first two properties in Proposition 3.23 ensure that this way $Poss$ is defined on all subsets of W as desired. The third property then determines possibility of a union as $Poss(A \cup B) = \max(Poss(A), Poss(B))$. Further for all $A \subseteq W$ we additionally require that $Poss(A) = 1$ iff $Pl(A) = \top$ and $Poss(A) = 0$ iff $Pl(A) = \bot$. Notice that it must be the case that $\max_{w \in W} Poss(w) = 1$. Hence, the resulting plausibility measure $Poss$ is in fact a possibility measure.

The fourth property ensures that we can update the cps with every plausible public announcement. It shall become clearer below. □

Let us define an analogue to default-equivalence between two plausibility spaces.

Definition 3.30. We say that a conditional plausibility space (W, Pl) (resp. conditional plausibility structure (W, Pl, π)) and a possibility space[25] $(W, Poss)$ (resp. a possibility structure $(W, Poss, \pi)$) are *default-equivalent* if for all disjoint subsets A and B of W it holds $Pl(A) <_D Pl(B)$ if and only if $Poss(A) < Poss(B)$.

[25] We use this notion to refer to the corresponding space for a possibility measure $Poss$.

Clearly, if the structures (W, Pl, π) and $(W, Poss, \pi)$ are default-equivalent, then they satisfy the same conditional beliefs.

Proposition 3.31. *Any conditional plausibility space* $(W, \mathcal{F}, \mathcal{F}', Pl)$ *satisfying the conditions in Proposition 3.23 is default-equivalent to the recovered possibility space* $(W, 2^W, Poss)$.

Proof. We want to show that for all disjoint subsets A and B of W, we have $Pl(A) <_D Pl(B)$ if and only if $Poss(A) < Poss(B)$. However, their default-equivalence follows directly from the construction of $Poss$. □

The idea of the following theorem is that if a cps is default-equivalent to a possibility space and they are updated with plausible public announcement $!P$, then the updated versions of these spaces are also default-equivalent. Whereas the cps is updated according to Definition 3.2, the possibility measure follows (A, α)-conditionalisation stated above and results in a new space $(W_P, 2^{W_P}, Poss_P)$ restricted to the new domain of the p-worlds. What is more, the desired properties of the cps are also preserved.

Theorem 3.32. *Any conditional plausibility space* $(W, \mathcal{F}, \mathcal{F}', Pl)$ *satisfying the conditions in Proposition 3.23 is such that the following property holds:*
If $Pl(P) \neq \perp$ *and the cps is default-equivalent to a possibility space* $(W, 2^W, Poss)$ *then the new cps* $(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P)$ *is default-equivalent to the possibility space obtained by updating* $(W, 2^W, Poss)$ *with public announcement* $!P$*. Moreover, the new cps also satisfies the conditions in Proposition 3.23.*

Proof. As expected we follow the path of the previous proofs for a single-agent plausibility space in Theorem 3.16 and a ranking function in Theorem 3.26. Therefore, we can update the standard cps with the function R_{PA} of public announcement in Definition 3.2.

Further, the possibility measure is modified by (A, α)-conditionalisation as described above in (3.2): $Poss_P(w) = Poss(w)/Poss(P)$ for all $w \in P$. Notice that it cannot be the case that $Poss(P) = 0$, since the original spaces are default-equivalent and it is given that $Pl(P) \neq \perp$.

In Theorem 3.16 we have also proved that the new cps satisfies the conditions in Proposition 3.23. Hence, the only thing left to show is that the both spaces remain default-equivalent after the update with public announcement $!P$, i.e., for all disjoint subsets A and B of W_P we have $Pl_P(A) <_D Pl_P(B)$ iff $Poss_P(A) < Poss_P(B)$. Then, indeed, their corresponding structures satisfy the same conditional beliefs:

(i) A possibility structure:
$(W_P, 2^{W_P}, Poss_P, \pi_P) \models B^{\psi}\varphi$ iff $Poss_P([\![\psi]\!]) = 0$ or $Poss_P([\![\varphi \wedge \psi]\!]) > Poss_P([\![\neg\varphi \wedge \psi]\!])$.[26]

(ii) A conditional plausibility structure:
$(W_P, \mathcal{F}_P, \mathcal{F}'_P, Pl_P, \pi_P) \models B^{\psi}\varphi$ iff $Pl_P([\![\psi]\!]) = \bot$ or $Pl_P([\![\varphi \wedge \psi]\!]) >_D Pl_P([\![\neg\varphi \wedge \psi]\!])$.

Recall that for both spaces the ordering on plausibilities (resp. possibilities) of sets of worlds is determined by the ordering on plausibilities (resp. possibilities) of worlds. Public announcement $!P$ eliminates the $\neg p$-worlds and what is more, in the case of possibility measures it also shifts all the 'surviving' p-worlds upwards towards 1 (expect for those with possibility 0). However, the original orderings on the p-worlds remain the same in both spaces. And as such both updates by public announcement preserve default-equivalence between the spaces, in particular, between their structures (see the summarising picture below).

It is rather straightforward that these updated structures satisfy the same formulas in $\mathcal{L}_\Phi$, since the propositional cases are obvious and we have just proved the case of conditional belief.

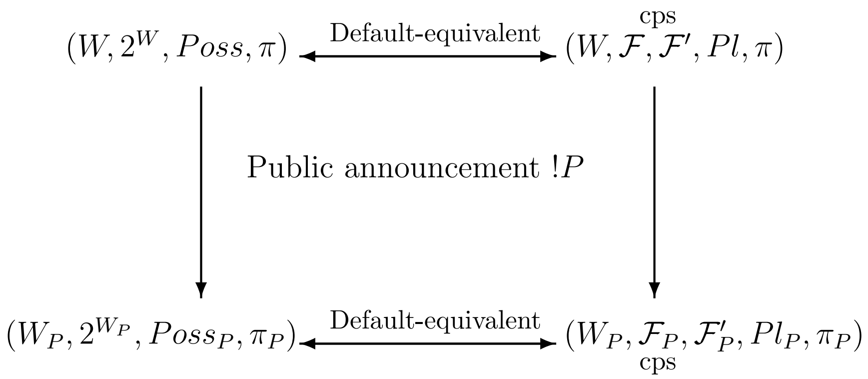

□

[26] We have relativised the rest of the possibility space to the p-worlds denoting it as W_P, 2^{W_P}, and π_P, respectively.

3.5 Embedded Update Mechanisms

3.5.1 Simulation of Public Announcement

In the previous sections we have shown that having a conditional plausibility space $(W, \mathcal{F}, \mathcal{F}', Pl)$ defined by CPl1–CPl5 (see Definition 2.7) with certain additional properties, we can recover a single-agent plausibility space $(W, \leq)$, κ-ranking $(W, 2^W, \kappa)$, and possibility space $(W, 2^W, Poss)$, respectively. Indeed, the added properties are far from random. Let us discuss them one by one:

1. $\mathcal{F} = 2^W$.
   All three frameworks stated above are defined on worlds. After all, their update mechanisms work on worlds as they take for their input a pair of world and action.

2. $D_W = \{Pl(V) : V \in \mathcal{F}\}$ is totally ordered by $\leq_D$: either $Pl(A) \leq_D Pl(B)$ or $Pl(B) \leq_D Pl(A)$ for all sets $A, B \in \mathcal{F}$. What is more, by CPl5 we have that for all $U \in \mathcal{F}'$, Pl_U places a total order on sets of $\mathcal{F}$.

3. $Pl(A \cup B) = \max(Pl(A), Pl(B))$ for all sets $A, B \in \mathcal{F}$.

   In [FrH01] an unconditional plausibility measure Pl is defined to be a ranking, if it satisfies the second and third property listed here (see Definition 1.21). Therefore, it is easy to verify that for all $U \in \mathcal{F}'$, Pl_U is a ranking.[27]

   Having all three properties, i.e., a ranking on 2^W, the plausibility of a set A is determined by the plausibilities of the worlds $w \in A$.[28]

4. The given cps $(W, \mathcal{F}, \mathcal{F}', Pl)$ is standard.
   This property enables conditioning for an arbitrary plausible public announcement as defined by R_{PA} in Definition 3.2.

5. $Pl(A) = \perp$ iff $A = \emptyset$.
   This property has its roots a bit deeper in the philosophical ground. As we

[27] Not to confuse this notion with a ranking function κ. Indeed, the latter is an instance of the former.

[28] Recall that in the case of a saps we have defined an ordering on sets in this fashion.

have already discussed in Section 2.2 on probabilistic conditioning, conditioning on sets with plausibility $\perp$, i.e., on implausible sets, is undesirable. Dealing with conditional belief $B^{\psi}\varphi \in \mathcal{L}_{\Phi}$, both ranking functions and possibility measures take the case when ψ is satisfied only in implausible (resp. impossible) worlds (i.e., $[\![\psi]\!] = \infty$, and 0 respectively) to be trivial resulting in $B^{\psi}\varphi$ being true for an arbitrary φ. However, a single-agent plausibility space takes the minimal elements of a given set, which means that if $[\![\psi]\!]$ consists only of 'implausible' (i.e., the maximal) worlds, they are (all) considered as the minimal. The only trivial case for a saps is $[\![\psi]\!] = \emptyset$, then indeed $B^{\psi}\varphi$ is true for an arbitrary φ. This is the reason for the extra fifth property to be required in the case of a saps.
However, from a practical point of view, we could allow the worlds to have as little plausibility as needed, but saving the value $\perp$ (∞, and 0 respectively) of implausibility only for the empty set. One could simply argue that as far as beliefs are concerned these irreversibly implausible worlds are rather dispensable anyway.

Corollary 3.33. *Let $(W, 2^W, \mathcal{F}', Pl)$ be a standard conditional plausibility space, where $Pl(.\ |W)$ is a ranking, and Pl^* is a ranking function κ or a possibility measure $Poss$. Then the following property holds:*
If $Pl(P) \neq \perp$ and the cps is default-equivalent to $(W, 2^W, Pl^*)$, then the new cps $(W_P, 2^{W_P}, \mathcal{F}'_P, Pl_P)$ is default-equivalent to the plausibility space obtained by updating $(W, 2^W, Pl^*)$ with public announcement $!P$.

(*) If moreover $Pl(A) = \perp$ implies $A = \emptyset$, then the property above also holds when a single agent plausibility space $(W, \leq)$ is substituted for $(W, 2^W, Pl^*)$.

The new cps is standard and $Pl(.\ |P)$ is a ranking. In the case (*) additionally $Pl(A\ |P) = \perp$ implies $A \cap P = \emptyset$.

Proof. Straightforward from the theorems 3.16, 3.26, and 3.32 (see the summarising picture below). □

We have shown that public announcement realised by the corresponding update mechanisms on single-agent plausibility models, ranking structures, and possibility structures, respectively, can be embedded in a cps with certain additional properties (i.e., a ranking $Pl(.\ |W)$ on 2^W). While on all the structures of these three classes we apply their update mechanisms, the cps is in principle updated purely by conditioning. We need to restrict the rest of the space, so that the result is a cps on W_P, but the main idea is to simply condition on P after a (plausible) public announcement $!P$.

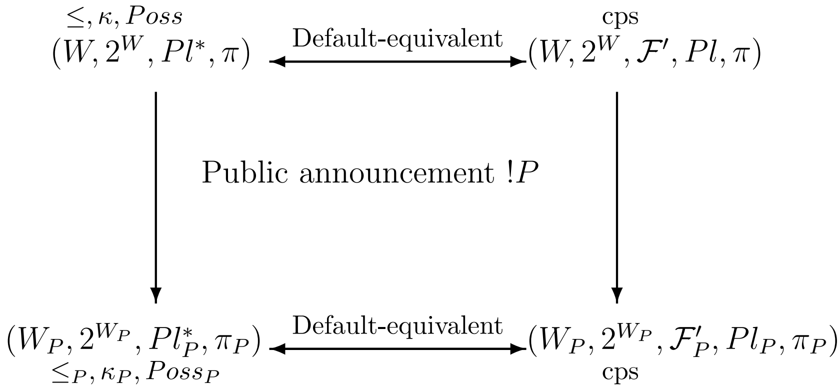

In a certain sense, we can simulate public announcement on single-agent plausibility models, ranking structures, and possibility structures, respectively, by conditioning on a conditional plausibility structure. At each step of the update process on the cps (e.g., in case there are several plausible public announcements in a row), we can construct a structure of one of these three classes which will satisfy the same formulas as it would if we had a structure of this class satisfying the same formulas as the original cps and applied its corresponding update mechanism all the way from the beginning.

3.5.2 Simulation of Radical Revision

Recall, that radical revision (resp. lexicographic revision) results in keeping the same states in W, but promoting all the p-worlds to be more plausible than all the $\neg p$-worlds, and within the two clusters the order remains unchanged. Since in our selected structures the order on plausibilities (resp. possibilities) of sets is determined by the order on plausibilities (resp. possibilities) of worlds, again default-equivalence between the structures will be preserved.

The idea for the update mechanisms would be to take an action model $\mathbf{\Sigma}$ with two plausible actions, one with precondition P and the other one with precondition $\overline{P}$. This way, all the worlds stay in model, but we can shift the two clusters relatively to each other. Obviously we shift the p-worlds upwards towards the top. But how much do we need to shift the $\neg p$-worlds to ensure that the two clusters are kept separated as required? In other words, what plausibility should we assign to the latter action to get the results of radical revision?[29]

Notice that the update mechanism for single-agent plausibility models (with anti-lexicographic ordering) does not succumb to this challenge and with its qualitative nature accounts for this kind of revision. On the other side, the quantitative update mechanisms for ranking structures, and possibility structures, respectively, do not necessarily behave in accordance with radical revision. Based on Spohn's (A, α)-conditionalisation, in each case we would have to fix an 'appropriate penalty', i.e., a sufficient α for $\overline{A}$, to ensure that the best worlds in $\overline{A}$ are still worse than the worst worlds in A.

In order to aim for some general results on this matter, in the following chapter we present an update model using plausibility measures.

[29] Cf. Jeffrey's Rule in Subsection 2.6.2.

Chapter 4

Extra Algebraic Properties

In the previous chapter we have seen how public announcement on single-agent plausibility models, ranking structures, and possibility structures, respectively, can be embedded in a conditional plausibility space. However, in order to account for different kinds of revision (e.g., radical or conservative) we need to define an actual update mechanism using cpm's. The update mechanism for Baltag and Smets' single-agent plausibility spaces is qualitative and as such uses the orderings. On the other hand, the update mechanisms for κ-rankings and possibility spaces are quantitative and naturally they make use of the present algebraic properties. Let us illustrate some kind of a generalisation of such a quantitative approach to updating.

4.1 General Model for Revision using Plausibility Measures

In this section we briefly present a revision model for epistemic state change using plausibility measures as suggested in [MaL08]. This model is general enough to subsume the conditionalisation of ranking functions (see Subsection 3.3.1), Jeffrey's Rule of probability updating (see Subsection 2.6.2), and the revision operator (3.2) in possibility theory.[1]

To make the subsequent discussion easier, we have the following. Let A be any set, for any binary relation $\leq$ over $A \times A$, $<$ is defined as $a < b$ iff $a \leq b$ and $b \nleq a$, and $=$ is defined as $a = b$ iff $a \leq b$ and $b \leq a$, for $a, b \in A$.

First, we need to define some simple and rational properties for operator $\otimes$.

Definition 4.1. Let $S = (W, \mathcal{F}, D, Pl)$ be a plausibility space, a, b, c be any elements in D and $\otimes$ be a mapping $D \times D \to D$, then $\otimes$ is called

[1] Note that this topic is beyond the scope of this thesis, but we consider it very relevant and include the main results (without proofs) of J. Ma and W. Liu for future research in this area. Understandably, this chapter is rather sketchy.

- *reversible* iff there exists a mapping $\otimes^{-1}$ such that $a \otimes^{-1} b \otimes b = a$ and $a \otimes b \otimes^{-1} b = a$ for $b \neq \bot$.
- *commutative* iff $a \otimes b = b \otimes a$.
- *associative* iff $a \otimes (b \otimes c) = a \otimes b \otimes c$.
- *equal-ranking* iff $(a \otimes b) \otimes^{-1} c = (a \otimes^{-1} c) \otimes b$ for $c \neq \bot$.
- *right-sign-keeping* iff $a \otimes c <_D b \otimes c$ for $a <_D b$.
- *left-sign-keeping* iff $c \otimes a <_D c \otimes b$ for $a <_D b$.
- *sign-keeping iff* $\otimes$ is both right-sign-keeping and left-sign-keeping.

Property equal-ranking says that an operation $\otimes$ and its reversing operation $\otimes^{-1}$ have the same level of operation grade, such as, '+' and its reverse '−' have the same level of arithmetic calculation grade and they are a grade lower than '×' and '/'.

Note that if $\otimes$ is reversible, then by setting $U_3 = W$ in Alg2 (see Definition 2.9), we obtain a conditional plausibility as follows

$$Pl(U_1 \,|U_2) = Pl(U_1 \cap U_2) \otimes^{-1} Pl(U_2).$$

The reason we need to have both the right-sign-keeping and left-sign-keeping conditions is that some operators may not be associative, so these two conditions are not totally equivalent.

Proposition 4.2. *Let $S = (W, \mathcal{F}, D, Pl)$ be a plausibility space and $\otimes$ be a reversible and right-sign-keeping mapping $D \times D \to D$, then $\otimes^{-1}$ is right-sign-keeping.*

Note that if $\otimes$ is commutative, then $\otimes$ is right-sign-keeping iff $\otimes$ is left-sign-keeping. But we still differentiate the two situations as there may be non-commutative operators, e.g., $\otimes^{-1}$.

Definition 4.3. Let $S = (W, \mathcal{F}, D, Pl)$ be a plausibility space and $\otimes$ be a mapping $D \times D \to D$, then $\otimes$ is called a *rational mapping* iff it satisfies reversible, commutative, associative, equal-ranking, and sign-keeping.

Proposition 4.4. *Let $S = (W, \mathcal{F}, D, Pl)$ be a plausibility space and $\otimes$ be a rational mapping $D \times D \to D$, then for any $a, b, c, d \in D$ and $b, c \neq \bot$, we have*

(i) $a \otimes^{-1} b \otimes^{-1} c = a \otimes^{-1} c \otimes^{-1} b$.

(ii) $a \otimes (d \otimes^{-1} c) = a \otimes d \otimes^{-1} c$.

(iii) $b \otimes^{-1} b = \top$.

In fact, when probability functions, ranking functions and possibility functions, are viewed as plausibility functions, the corresponding $\otimes$s (which are '$\times$', '$+$', and '$\times$' respectively) are indeed rational mappings.

We define the revision model by plausibility measures as follows.

Definition 4.5. Let $S = (W, 2^W, D, Pl)$ be a plausibility space for the prior state, and $S_e = (W, \mathcal{F}_e, D, Pl_e)$ be the plausibility space for new evidence where $\mathcal{F}_e = 2^{\{A_1,\ldots,A_n\}}$ is the powerset of a partition of W, then the *revised plausibility measure*, denoted as Pl_{re}, is

$$Pl_{re}(w) = Pl_e(A_i) \otimes^{-1} Pl(A_i) \otimes Pl(w),\ w \in A_i,\ 1 \leq i \leq n.$$

Proposition 4.6. *Let $S = (W, 2^W, D, Pl)$ be a plausibility space for the prior state, and $S_e = (W, \mathcal{F}_e, D, Pl_e)$ be the plausibility space for new evidence where $\mathcal{F}_e = 2^{\{A_1,\ldots,A_n\}}$, then we have*

$$Pl_{re}(A_i) = Pl_e(A_i),\ 1 \leq i \leq n.$$

This proposition shows that the above definition indeed preserves the value $Pl_e(A_i)$ from the evidence, so it satisfies the general requirement in revision that the new evidence has to be preserved.

Here are some general properties of the revision by plausibility measures.

Proposition 4.7. *Let $S = (W, 2^W, D, Pl)$ be a plausibility space for the prior state and $S_{e1} = (W, \mathcal{F}_{e1}, D, Pl_{e1})$, $S_{e2} = (W, \mathcal{F}_{e2}, D, Pl_{e2})$ be two plausibility spaces for two new pieces of evidence such that $\mathcal{F}_{e1} = \mathcal{F}_{e2} = 2^{\{A_1,\ldots,A_n\}}$, then we have $(Pl_{re1})_{re2} = Pl_{re2}$.*

This proposition reveals that if two pieces of evidence are about the same event but differ on the strengths, then the evidence arriving last will suppress the former (cf. Jeffrey's Rule in Subsection 2.6.2).

When new evidence is given on $\mathcal{F}_e = 2^{\{A,\overline{A}\}}$ within a plausibility space, the above revision is reduced to the well known (A, α)-conditionalisation of ranking functions ([Spo88], [DuP93]), which is the revision when $S_e = (W, 2^{\{A,\overline{A}\}}, D, Pl_e)$ such that $Pl_e(A) = \top$ and $Pl_e(\overline{A}) = \alpha$. Thus we have

$$Pl_{A,\alpha}(w) = \begin{cases} Pl(w) \otimes^{-1} Pl(A) & \text{for } w \in A \\ \alpha \otimes^{-1} Pl(\overline{A}) \otimes Pl(w) & \text{for } w \in \overline{A}. \end{cases}$$

Definition 4.8. Let $\oplus$ be a mapping $D \times D \to D$, then $\oplus$ is called *bounded-additive* if and only if it follows: $\top \oplus d = d \oplus \top = \top$ for all $d \in D$.

For convenience, if Pl is associated with a bounded-additive $\oplus$, then we simply call Pl bounded-additive. It is clear to see that ranking functions and possibility measures are bounded-additive, but unfortunately, the probability measure is not bounded-additive.

The revision by Definition 4.5 (using Proposition 4.6) can be equivalently rewritten as

$$Pl_{re}(w) \otimes^{-1} Pl(w) = Pl_{re}(A_i) \otimes^{-1} Pl(A_i),\ w \in A_i,\ 1 \leq i \leq n.$$

It is a counterpart of so called probability kinematics in probability theory.[2] It has been proved that Jeffrey's Rule follows probability kinematics.[3] Hence the revision strategy described here can be called plausibility kinematics.

We give the formal definition of plausibility kinematics as follows.

Definition 4.9. Suppose that two plausibility measures Pl and Pl^* disagree on the plausibility values they assign to a set of mutually exclusive and exhaustive events $A_1, \ldots, A_n$. The distribution Pl^* is said to be obtained from Pl by *plausibility kinematics* on $A_1, \ldots, A_n$, if and only if for any $w \in A_i$, $1 \leq i \leq n$, it holds

$$Pl^*(w) \otimes^{-1} Pl(w) = Pl^*(A_i) \otimes^{-1} Pl(A_i).$$

[2] Originally from Jeffrey, R. C. *The Logic of Decision.* New York: McGraw-Hill, 1965. 2nd ed. Chicago: University of Chicago Press, 1983.

[3] In Chan, H., and A. Darwiche. "On the Revision of Probabilistic Beliefs using Uncertain Evidence." *Artificial Intelligence* 163 (2005): 67–90.

Obviously, the revision strategy in Definition 4.5 shows that the revised plausibility measure is obtained from the prior plausibility measure by plausibility kinematics.

Next it is proved that the revision strategy does achieve a minimal change. Namely, we show that among all revision strategies, the plausibility measure obtained by plausibility kinematics has the shortest distance to the prior plausibility measure.

First, we define a distance function, which is generalised from its probability counterpart (ibid. footnote 3).

Definition 4.10. Let Pl and Pl^* be two plausibility measures on 2^W, then the *distance* between Pl and Pl^* is defined as

$$d(Pl, Pl^*) = \odot(max_w Pl^*(w) \otimes^{-1} Pl(w)) - \odot(min_w Pl^*(w) \otimes^{-1} Pl(w)),$$

where we define $\bot \otimes^{-1} \bot = \top$, and $\odot$ is a mapping $D \to R$ and satisfies the following:

1. $\odot(a \otimes^{-1} b) = \odot a - \odot b$.
2. if $a < b$, then $\odot a < \odot b$.
3. $\odot\bot = \infty$.

Pl and Pl^* are said to have the same support if for all w it holds $Pl(w) \neq \bot$ iff $Pl^*(w) \neq \bot$. If Pl and Pl^* do not have the same support, as $\odot\bot = \infty$, we can conclude that $d(Pl, Pl^*) = \infty$.

Proposition 4.11. *$d(Pl, Pl^*)$ defined in Definition 4.10 is a distance function.*

A common perspective on revision strategies is to have a minimal change between the prior belief (resp. epistemic state) and the revised belief (resp. epistemic state). The theorem below shows that the suggested revision strategy is optimal in the sense that it satisfies this common perspective.

Theorem 4.12. *The plausibility distribution Pl^* obtained from Pl by plausibility kinematics on partition $A_1, \ldots, A_n$ of W is optimal in the following sense. Among all possible plausibility distributions that agree with Pl on the plausibility values of events $A_1, \ldots, A_n$, Pl^* is the closest to Pl according to the distance measure by Definition 4.10.*

Conclusion

At the beginning we have introduced different kinds of representations of uncertainty. Then, we have shown how to update each of these frameworks. In both cases we have focused on a (conditional) plausibility space, since it generalises all the other representations. We have also presented algebraic properties and Jeffrey's Rule, both based on the ideas from probability theory.

The third chapter is the core of this thesis. We have selected three frameworks and for each of them listed the properties that need to be added to a conditional plausibility space in order to recover these frameworks. We have proved that public announcement on single-agent plausibility models (in [BaS08]), ranking structures, and possibility structures, respectively, realised by their corresponding update mechanisms can be embedded in the framework of conditional plausibility spaces. It has turned out that the focus falls on a satisfaction condition for conditional beliefs. This has lead to the definition of default-equivalence between a structure of one of the classes stated above (resp. its space) and a conditional plausibility structure (resp. space). At the end of the chapter we have also discussed radical revision on the selected frameworks and its potential to be embedded in a conditional plausibility space in similar way as public announcement. However, the update mechanisms of quantitative nature do not necessarily behave according to radical revision unless an 'appropriate penalty' is fixed in each case. In order to investigate this matter further we have aimed for some kind of a generalisation of quantitative update mechanisms.

In the fourth chapter we have briefly presented a revision model for epistemic state change using plausibility measures as suggested in [MaL08]. It has similar features as the quantitative update mechanism for ranking structures, but it deals only with mutually exclusive and exhaustive events (i.e., on a partition of the set of worlds W). After all, this revision model has been created in the spirit of Jeffrey's Rule and Spohn's (A, α)-conditionalisation. This topic, while beyond the scope of this thesis, is considered very relevant and provides grounds for future research.

"Conditioning is a wonderful tool…"[Hal03, p. 74] with great potential. A little bit more polishing and the dark corners of uncertainty shall become its light…

Bibliography

[ABD13] Andersen, M. B., T. Bolander, H. van Ditmarsch, and M. H. Jensen. "Bisimulation for Single-Agent Plausibility Models." *AI 2013: Advances in Artificial Intelligence.* Eds. S. Cranefield, and A. Nayak. Dordrecht: Springer, 2013. 277–288.

[Auc03] Aucher, Guillaume. A Combined System for Update Logic and Belief Revision. *Master's Thesis.* ILLC University of Amsterdam, 2003.

[Auc07] Aucher, Guillaume. "Interpreting an Action from What We Perceive and What We Expect." *Journal of Applied Non-Classical Logics* 17 (2007): 9–38.

[BaS08] Baltag, A., and S. Smets. "A Qualitative Theory of Dynamic Interactive Belief Revision." *Logic and the Foundations of Game and Decision Theory (LOFT7).* Eds. G. Bonanno, W. van der Hoek, and M. Wooldridge. Amsterdam: Amsterdam University Press, 2008. 11–58. Vol. 3 of *Texts in Logic and Games.*

[Ben04] van Benthem, Johan. "Dynamic Logic for Belief Revision." *Journal of Applied Non-Classical Logics* 17.2 (2007): 129–155.

[DHK06] van Ditmarsch, Hans, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic.* Dordrecht: Springer, 2008.

[DuP93] Dubois D., and H. Prade. "Belief Revision and Updates in Numerical Formalisms: An Overview, with New Results for the Possibilistic Framework." *Procs. of IJCAI* (1993): 620–625.

[FrH95] Friedman, Nir, and Joseph Y. Halpern. "Plausibility Measures: A User's Guide." *Conference on Uncertainty in Artificial Intelligence* (1995): 175–184.

[FrH97] Friedman, Nir, and Joseph Y. Halpern. "Modeling Belief in Dynamic Systems, Part I: Foundations." *Artificial Intelligence* 95 (1997): 257–316.

[FrH01] Friedman, Nir, and Joseph Y. Halpern. "Plausibility Measures and Default Reasoning." *JACM* 48.4 (July 2001): 648–685.

[GHJ14] Gierasimczuk, Nina, Vincent F. Hendricks, and Dick de Jongh. "Logic and Learning." *Johan van Benthem on Logic and Information Dynamics.* Eds. A. Baltag, and S. Smets. Dordrecht: Springer, 2014. 267–288. Vol. 5 of *Outstanding Contributions to Logic.*

[Hal03] Halpern, Joseph Y. *Reasoning about Uncertainty.* Cambridge, MA: MIT Press, 2003.

[MaL08] Ma, Jianbing, and Weiru Liu. "A General Model for Epistemic State Revision using Plausibility Measures." *Procs. of ECAI* (2008): 356–360.

[Spo88] Spohn, W. "Ordinal Conditional Functions: A Dynamic Theory of Epistemic States." *Causation in Decision, Belief Change, and Statistics.* Eds. W. L. Harper, and B. Skyrms. Dordrecht: Kluwer, 1988. 105–134. Vol. 2 of *Procs. of the Irvine Conference on Probability and Causation.*