

---

# From conditional probability to the logic of doxastic actions

---

Alexandru Baltag\*  
Computing Laboratory  
Oxford University  
Oxford OX1 3QD, UK.

Sonja Smets†  
Center for Logic and Philosophy of Science‡  
Vrije Universiteit Brussel  
Brussels B1050, Belgium  
& Research Group on the Philosophy of Information  
Oxford University  
Oxford OX1 3QD, UK

## Abstract

We investigate the discrete (finite) case of the Popper-Renyi theory of conditional probability, introducing *discrete conditional probabilistic models for knowledge and conditional belief*, and comparing them with the more standard plausibility models. We also consider a related notion, that of *safe belief*, which is a weak (non-negatively introspective) type of “knowledge”. We develop a probabilistic version of this concept (“degree of safety”) and we analyze its role in games. We completely axiomatize the logic of conditional belief, knowledge and safe belief over conditional probabilistic models. We develop a theory of probabilistic dynamic belief revision, introducing “action models” and a notion of probabilistic update product, that comes together with appropriate reduction laws.

## 1 Introduction

In this paper, we connect three commonly-used approaches for representing beliefs and belief-updates: 1) the *Bayesian approach in its extended Popper-Renyi (-de Finetti) version* (based on conditional probabilistic models allowing conditionalization on events possibly having zero probability), 2) the *classical AGM-style Belief Revision approach* in its *semantic* presentation (based on plausibility models and plausibility ranking) and 3) the *modal DEL-style<sup>1</sup> approach* (based on the distinction between “static” belief revision, represented using conditional belief operators, and “dynamic” revision, represented using epistemic/doxastic action models, dynamic modalities and some notion of product update). Some connections between the first two were

---

\*E-mail: baltag@comlab.ox.ac.uk

†E-mail: sonsmets@vub.ac.be

‡ Sonja Smets’ post-doctoral research is sponsored by the Flemish Fund for Scientific Research.

<sup>1</sup>i.e. “Dynamic-Epistemic Logic”, in the tradition of [17, 7, 6, 5] and the work of J. van Benthem and the “Amsterdam school”.

already studied, most thoroughly in [15, 13, 19, 2], where it was shown that a correct probabilistic understanding of belief revision and conditional beliefs requires an extension of classical probability theory, along the lines of the Popper-Renyi axioms. The connections between the second and the third approach are investigated in a number of recent papers [3, 32, 28, 8, 10, 9, 11], while [27, 29, 21] relate the first and the third approach only briefly (using only classical probabilistic models).

Combining these approaches into one, we introduce a “*qualitative*” *dynamic logic of conditional beliefs, knowledge, safe belief and belief-updating actions*, which is *decidable and complete with respect to (Popper-style) conditional-probabilistic models*. The syntax and the proof system for this logic are the same we introduced in [10, 9, 11], but the semantics is probabilistic (instead of using plausibility models). We develop a theory of dynamic belief revision over probabilistic models, by introducing “action models” and a notion of *conditional-probabilistic product update*, which generalizes to a belief-revision context the corresponding notion introduced in [21, 29] for probabilistic epistemic actions. One can also extract from the probabilistic update product a corresponding qualitative update notion for plausibility models; this last notion coincides with the “*Action-Priority Update*”<sup>2</sup>, introduced in our previous work [10, 9, 11], as a way of combining the “update product” from [7, 8] with ideas from Belief Revision theory.

This paper assumes the general distinction, made in [32, 8, 28], between “*dynamic*” and “*static*” belief revision. To summarize it: “static” belief revision, corresponding closely to the classical AGM theory in [1, 16] and embodied in our setting by the conditional belief operators  $B_a^P Q$ , captures *the agent’s changing beliefs about an unchanging world*; but since in fact, in a modal logic setting, the world is *always changed by our changes of beliefs*, the best way to understand a doxastic conditional  $B_a^P Q$  is as saying that *after learning P, agent a believes that Q was the case (before*

---

<sup>2</sup>This name, used in [11], was proposed by J. van Benthem. In [10, 9], this was called “anti-lexicographic product update”.

*the learning*); in contrast, “dynamic” belief revision uses dynamic modalities to capture *the agent’s revised beliefs about the world as it is after revision*:  $[P!]B_a Q$  says that *after learning P, agent a believes that Q is the case (in the world after the learning)*. The standard alternative [20] to the AGM theory calls this *belief update*, but like the AGM approach, it only deals with “first-level” beliefs (about ontic facts) from a non-modal perspective, neglecting any higher-order “beliefs about beliefs”; as a result, it completely misses the changes induced in the world (including in the other agents’ epistemic states) by our belief-updating actions (e.g. the learning of a Moore sentence); this is shown by the acceptance in [20] of the AGM “Success Axiom”: in dynamic notation, the setting in [20] validates the axiom  $[P!]B_a P$  (which cannot accommodate Moore sentences); instead, the authors of [20] exclusively concentrate on the possible changes of (ontic) facts that may have occurred during our learning. In contrast, our approach to belief update (following the DEL tradition) may be thought of as “dual” to the one in [20]: we completely neglect here the ontic changes<sup>3</sup>, considering a world in which the only changes are induced by “*purely doxastic*” actions (such as learning, updates, discovery, communication etc.).

As in [9] (but now with a probabilistic semantics), we introduce a “weak” (non-negatively introspective) notion of knowledge, called *safe belief* and corresponding to Stalnaker’s concept of “knowledge as defeasible belief” in [26]. But in the context of probabilistic models, we refine this notion introducing a quantitative scale of “*degrees of safety*” of a given belief, giving examples from Game Theory to illustrate the usefulness of this notion.

All our completeness results are original work, even if the proof systems are known: it is a highly non-trivial fact that the same logics are complete for the (*discrete, finite*) conditional probabilistic models.<sup>4</sup>

## 2 Conditional Doxastic Logic over Probabilistic Models

It is well known that simple probability measures yield problems in the context of describing an agent’s beliefs and how they can be revised.

First, it seems natural to assume (and is usually assumed) that *beliefs are closed under finitary conjunctions* (and that moreover the belief operator satisfies Kripke’s axiom *K*). But this means that we cannot identify “(certain) belief” with probability  $\geq p$  (“high degree of certainty”) for any  $p < 1$  (since none of these is closed under finitary conjunc-

<sup>3</sup>But our approach can be easily modified to incorporate ontic changes, along the lines of [31].

<sup>4</sup>The non-trivial character of this fact is also witnessed by our need to limit ourselves to the finite discrete case. The problem of finding a complete logic for the general (infinite) case is still open.

tions). So “(certain) belief” must imply “(subjective) probability 1”. This point can be dramatized into the so-called Lottery Paradox. As put in [15], in a lottery composed of 1000 tickets which is believed to be fair, an agent assigns probability 0.99 to the belief that any particular ticket is not the winning one; on the other hand, he obviously will not assign probability 0.99 to the belief that *no* ticket is the winning one! On the contrary, he should assign probability 0 to this belief: some ticket will definitely be winning.

The received wisdom in Bayesianism is that *learning new information corresponds to probabilistic conditionalization*. But once “belief” is accepted to imply “subjective probability 1”, any non-trivial belief revision (triggered by learning that one of the agent’s beliefs was false) will correspond to *conditioning on events of measure 0*: an impossible task in the classical (Kolmogorovian) theory, in which conditional probability is defined as a ratio (whose denominator will be 0 in this case). In probabilistic applications, e.g. in Game Theory, this problem is sometimes preempted by requiring that only impossible events are assigned probability 0. But this, in effect, is a way of eluding the problem by simply stipulating that *agents never have any wrong beliefs*. In fact, this *collapses belief into knowledge*: in a finite discrete probabilistic space<sup>5</sup> satisfying this stipulation<sup>6</sup>, the “belief” operator becomes equivalent to a quantifier over all the states of the space; but this is the standard definition of *knowledge*! The unavoidable conclusion is that Bayesian belief update, based on standard Probability Theory, simply cannot deal with any non-trivial belief revision.

There are several possible solutions to this problem. In this paper we adopt the Popper-Renyi theory of conditional probabilities [23, 14, 15, 25, 24, 18], which takes conditional probability as basic instead of simple probability, and which was already applied to belief revision in [15, 2, 19]. We focus here on the *discrete finite* case, which allows us to simplify the definitions (and provide a unique, canonical setting<sup>7</sup> and to easier compare this with the qualitative (plausibility-based) settings for belief revision. A *discrete conditional probability space (dcps, for short)* is a pair  $(S, \mu)$ , where  $S$  is a *finite* set of states and  $\mu : \mathcal{P}(S) \times \mathcal{P}(S) \rightarrow [0, 1]$  satisfies the following axioms:

1.  $\mu(A|A) = 1$ ,
2.  $\mu(A \cup B|C) = \mu(A|C) + \mu(B|C)$ , if  $A \cap B = \emptyset$ ,  $C \neq \emptyset$ ,
3.  $\mu(A \cap B|C) = \mu(A|B \cap C) \cdot \mu(B|C)$ .

<sup>5</sup>I.e. a finite state space such that all its subsets are measurable.

<sup>6</sup>In the finite discrete case, this says that every state has a non-zero probability.

<sup>7</sup>The various axiomatic settings proposed in [23, 25, 14] become equivalent in the discrete case, and moreover they are equivalent in this case with the *conditional lexicographic probability spaces* proposed by researches in Game Theory.

In fact, all the information about  $\mu$  is captured by probabilities conditioned only on *pairs of states*, i.e. by the quantities:

$$(s, t)_\mu := \mu(s|\{s, t\})$$

for all  $s, t \in S$ . We skip the subscript when the measure is understood, and we call  $(s, t)$  the *priority degree*, or the *degree of superiority*, of  $s$  with respect to  $t$ . This gives us an *alternative description* of discrete conditional probability spaces, as *priority spaces*: binary probabilistic assignments  $(\bullet, \bullet) : S \times S \rightarrow [0, 1]$  on a finite set  $S$ , satisfying the axioms:

$$\begin{aligned} (s, s) &= 1; \\ (t, s) &= 1 - (s, t), \text{ for } s \neq t; \\ (s, w) &= \frac{(s, t) \cdot (t, w)}{(s, t) \cdot (t, w) + (w, t) \cdot (t, s)} \end{aligned}$$

for  $s \neq w$  and denominator  $\neq 0$ . It is easy to see that *any priority space*  $(S, (\bullet, \bullet))$  *uniquely determines a dcps*  $(S, \mu)$  such that  $(s, t) = (s, t)_\mu$  for all  $s, t \in S$ .

The (*non-strict*) *priority relation*  $\leq \subseteq S \times S$ , defined on states by putting  $s \leq t$  iff  $(s, t) \neq 0$ , is a special case of (the converse of) the “superiority” relation introduced in [15] (following De Finetti). In fact, the *strict* priority relation, which can be easily seen to be given by

$$s < t \text{ iff } (t, s) = 0 \text{ iff } (s, t) = 1, s \neq t,$$

was introduced in [2] under the name of “ranking ordering”. It is easy to see that (in the finite discrete case to which we are confined here) *the priority relation is a total preorder*<sup>8</sup>. By arranging the states of a finite space of size  $n$  in a list of non-decreasing priority  $s_1 \leq s_2 \leq \dots \leq s_n$ , we see that *in order to specify a discrete space of size  $n$  it is enough to give  $n - 1$  independent (non-zero) conditional probabilities, namely the priority degrees  $(s_i, s_{i+1})$  for  $1 \leq i \leq n - 1$ .*

In a dcps, there is a straightforward way to define (conditional) belief and knowledge, by simply identifying “belief” with “probability 1”<sup>9</sup> and “knowledge” with “true in all states”. In other words, for a “proposition”  $P \subseteq S$ , we put:  $B^P Q$  iff  $\mu(Q|P) = 1$ , and  $KP$  iff  $P = S$ . “Belief” is defined as “belief conditional on a tautology”, i.e. we put  $BP := B^S P$ . But these definitions assume that the state space  $S$  is restricted to all the states that the (implicit) agent considers as *epistemically possible*. In a more general context (as the multi-agent case considered below), knowledge has to be defined by quantifying only over epistemically possible states, while in the definition of belief we have to conditionalize  $\mu$  on the set of epistemically possible states.

<sup>8</sup>[2] extends this to countably additive probability measures.

<sup>9</sup>Van Fraassen [15] and Arlo-Costa and Parikh [2] consider more subtle distinctions (such as “full belief” and “plain belief”), but it is easy to see that in the case of discrete spaces all these notions become equivalent with “belief” as defined here.

Given a (finite) set  $\mathcal{A}$  of “agents”, a *discrete conditional probability frame* (or *dcpf*, for short) is a structure  $(S, \mu_a, \Pi_a)_{a \in \mathcal{A}}$ , such that, for each  $a \in \mathcal{A}$ ,  $(S, \mu_a)$  is a discrete conditional probability space and  $\Pi_a$  is a partition (the “information partition”) of  $S$ . Equivalently, we can of course use equivalence relations  $\sim_a$  instead of partitions  $\Pi_a$ . For a state  $s$ , denote by  $s(a)$  the *information cell* of  $s$  in the partition  $\Pi_a$  (or the  $\sim_a$ -*equivalence class* of  $s$ ). *Knowledge* and (*conditional*) *belief* become now dependent on the (information the agent possesses about the) state:

$$B_a^P Q := \{s \in S : \mu(Q|P \cap s(a)) = 1\},$$

$$K_a P := \{s \in S : s(a) \subseteq P\}.$$

We interpret the conditional belief statement  $s \in B_a^P Q$  in the following way: if the actual state is  $s$ , then after “learning” that  $P$  is the case (in the state  $s$ ), agent  $a$  will believe that  $Q$  was the case (at the same state  $s$ , i.e. before the learning). We abbreviate  $(s, t)_{\mu_a}$  as  $(s, t)_a$ . We denote by  $\leq_a$  the induced priority relation. For a given set  $\Phi$  of *atomic sentences* (intuitively denoting “ontic facts” about the world), a *discrete conditional probability model* (*dcpm*, for short) is, as usually in modal logic, a structure  $\mathbf{S} = (S, \mu_a, \Pi_a, \|\bullet\|)$  consisting of a dcpf  $(S, \mu_a, \Pi_a)$  together with a valuation  $\|\bullet\| : \Phi \rightarrow \mathcal{P}(S)$ .


In fact, for the doxastic-epistemic logic (i.e. for computing  $K_a$ ,  $B_a$  and  $B_a^P$ ), *only the priority degrees between distinct, but epistemically indistinguishable states are relevant*. So it is enough to know  $\{(s, t)_a : s \sim_a t\}$ . One can thus consider the *local priority* relation  $\triangleleft_a$ , defined as the intersection of the relations  $\leq_a$  and  $\sim_a$ . We denote by  $\triangleleft_a$  the corresponding *strict* relation. It is easy to see that, for  $s \neq t$ , we have

$$s \triangleleft_a t \text{ iff } t \in B_a^{\{s, t\}} s.$$


In other words, *the (local) priority relation actually captures a notion of priority of beliefs*: given the information that the actual state is one of two different states  $s$  or  $t$ , agent  $a$  will believe it is  $s$  (with probability 1) iff  $s \triangleleft_a t$ . So in representing a dcpm, we will only give the priority degrees between successive distinct states in the same partition cell (listed in non-decreasing local priority order). This provides a way to encode the information partition itself into the probabilistic information, so we *do not* have to represent the indistinguishability relations as well.

**Example 1:** Alice and Bob play a game, in which an (anonymous) referee takes a coin and puts it on the table in front of them, lying face up but in such a way that the face is covered (so Alice and Bob cannot see it). The goal of the game is to guess which face is up. Based on previous experience, (it is common knowledge that) Alice and Bob believe that the upper face is Heads (-since e.g. they noticed that the referee had a strong preference for Heads). And in fact, they’re right: the coin lies Heads up. A model


for this situation can be represented as:



**Example 2:** Compare this with the case in which the agents are not completely certain, but in which (it is common knowledge that) they believe with a probability of 0.99 that the coin lies Heads up:



**Example 3:** The situation is as in Example 2, except that it is common knowledge that Bob *can see* the face of the coin. Not only his probabilities become trivial, but his information partition will consist of singletons, so we will not even have to mention his priority degrees: in each case, he knows the real state of the system.




## 2.1 Relating probabilistic and relational models of conditional belief

To compare probabilistic models with *qualitative (relational)* ones, we introduce Kripke models for knowledge and conditional belief based on plausibility relations. A *finite (epistemic-doxastic) plausibility frame*<sup>10</sup> is a structure  $(S, \leq_a, \sim_a)_a$ , where  $S$  is finite and, for each  $a$ ,  $\leq_a$  is a *total* (i.e. “connected”, or “complete”) preorder and  $\sim_a$  is an equivalence relation. A plausibility model is a plausibility frame together with a valuation. In a plausibility model, knowledge  $K_a$  is defined in the standard way (as a Kripke modality) using  $\sim_a$ , while conditional belief is defined as:

$$B_a^P Q := \{s \in S : \text{Min}_{\leq_a} P \cap s(a) \subseteq Q\}$$

where  $\text{Min}_{\leq_a} T = \{s \in T : s \leq_a t \text{ for all } t \in T\}$  and  $s(a) = \{t \in S : s \sim_a t\}$ .

Note that (as in the case of probability models) only the plausibility relation between states in the same information cell are relevant; in other words, only the “local” plausibility relation  $\trianglelefteq_a := \leq_a \cap \sim_a$  is needed. In fact, this relation encodes the epistemic relations  $\sim_a$  as well. So, as before, we only represent  $\trianglelefteq_a$ , and for convenience we skip all the loops (since  $\trianglelefteq$  is reflexive anyway). Now, Example 1 above becomes:



For both local and global plausibility relations, we can also consider their “strict” versions  $s <_a t$  and  $s \triangleleft_a t$ . Finally,

<sup>10</sup>The notion here is the one we introduced in [8, 10, 9], but it is closely related to other notions in the literature: “Grove models”, “Lewis spheres”, “Spohn ordinal plausibility ranking”.

the relation of “equi-plausibility” is the equivalence relation  $\cong_a$  induced by the preorder  $\trianglelefteq_a$ :  $s \cong_a t$  iff  $s \trianglelefteq_a t$  and  $t \trianglelefteq_a s$ .

One usually reads  $s \leq_a t$  as saying that *state s is “at least as plausible” as state t*. But it is easy to see that we cannot identify “more plausible” with “higher probability”: if we accept the identification of (plausibility-based) conditional belief with “(conditional) probability 1”, then we *cannot* have  $s \leq_a t$  iff  $\mu_a(s) \geq \mu_a(t)$ . For the same reason, the plausibility preorder is not even given by the order of *relative (conditional) degrees of belief*: we *cannot* have  $s \leq_a t$  iff  $(s, t)_a \geq (t, s)_a$ . The most natural (global) plausibility relation  $\leq_a$  giving rise to the same notion of belief as a given dcpm  $\mathbf{S}$  is the *priority* relation ( $s \leq_a t$  iff  $(s, t)_a \neq 0$ ). In fact, it is easy to see that *the only local plausibility relation  $\trianglelefteq_a$  on a given dcpm  $\mathbf{S}$  that gives rise to the same (conditional) belief operator (as the one defined probabilistically) is the local priority relation*. This means that “plausibility” is in fact a *qualitative* notion, expressing (not degrees, or intensity, of beliefs, but) “*priority*” of beliefs:  $s \triangleleft_a t$  means simply the same as  $s \in B_a^{\{s,t\}} s$ . The plausibility relations can only capture “firm” (though conditional) beliefs, i.e. *believed with (conditional) probability 1*.

*So any discrete conditional probability model is a plausibility model (with the same notions of knowledge and conditional belief).*

**Lemma 1.** *Conversely, every finite plausibility model can be “probabilized”<sup>11</sup>: we can define conditional probability measures for each agent, that will give rise to the same conditional beliefs.*

This correspondence between the two types of models can be used to prove *completeness with respect to dcpm’s of the conditional doxastic logic CDL*, introduced in [8, 10, 9]:

**Theorem 1.** The logic *CDL* introduced in [8], and presented in Appendix A below, is *complete for dcpm’s*.

## 3 Safe Belief and Degrees of Safety

Stalnaker’s *defeasibility analysis* of knowledge in [26] is based on the idea that “if a person has knowledge, than that person’s justification must be sufficiently strong that it is not capable of being defeated by evidence that he does not possess” (Pappas and Swain [22]). However, Stalnaker interprets “evidence” as “true information”, and thus the concept of knowledge in [26] differs from ours. Unlike the one in [26], our notion of knowledge is *negatively introspective* and corresponds to interpreting “evidence” in the above quote as meaning “any information, be it truthful or not”. Thus, our “knowledge” is more robust than Stalnaker’s: it resists any belief revision, it is not capable of

<sup>11</sup>But this cannot generally be done in a unique way.

being defeated by *any* evidence (including false evidence). Nevertheless, we consider the weaker concept to be equally important, and so in [10, 9, 11] we introduced it (in the context of plausibility models) under the name of “safe belief”.

Since dcpm’s are plausibility models (with “priority” as the plausibility relation), we can use the same definition: the “safe belief” operator is the Kripke modality  $\Box_a$  associated to the converse  $\succeq_a$  of the local priority relation<sup>12</sup>, i.e. given by

$$\Box_a Q := [\succeq_a]Q = \{s \in S : \forall t \in S (t \preceq_a s \rightarrow t \in Q)\},$$


for all  $S$ -propositions  $Q \subseteq S$ . We read  $s \in \Box_a Q$  as saying that: *at state  $s$ , agent  $a$ ’s belief of  $Q$  (being the case) is safe*; or *at state  $s$ ,  $a$  safely believes that  $Q$* . An important observation is that this notion *does indeed capture Stalnaker’s non-standard concept of “knowledge”*:

$$s \in \Box_a Q \text{ iff } s \in B_a^P Q \text{ for all } P \subseteq S \text{ such that } s \in P.$$

So *safe beliefs are precisely the beliefs which are persistent under revision with any true information*. Another important observation, already made in [9], is that *one can characterize conditional belief only in terms of knowledge and safe belief*: if  $\tilde{K}_a P = \neg K_a \neg P$  is the Diamond modality for  $K$ , then

$$B_a^P Q = \tilde{K}_a P \rightarrow \tilde{K}_a (P \wedge \Box_a (P \rightarrow Q)).$$

**Example 4: Dangerous Learning.** This starts with the situation in Example 1. Alice has to get out of the room for a minute, which creates an opportunity for Bob to quickly raise the cover in her absence and take a peek at the coin. He does that and so he sees that the coin is Heads up. After Alice returns, she obviously doesn’t know whether or not Bob took a peek at the coin, but she believes he didn’t do it: taking a peek is against the rules of the game, and so she trusts Bob not to do that. The situation after this action is described by the following model  $S'$ , with actual state  $s'_1$  in the left upper corner:



In both Examples 1 and 4 above, Alice holds a *true belief* (at the real state) that the coin lies Heads up: the actual state satisfies  $B_a H$ . In both cases, this true belief is *not*

<sup>12</sup>A similar notion was defined in [30] in a different context, under the name of “preference modality”.

*knowledge* (since Alice doesn’t know the upper face); nevertheless, in Example 1, this belief is *safe* (although it is *not known by the agent to be safe*): no additional truthful information (about the real state  $s$ ) can force her to revise this belief. To see this, note that any *new* truthful information would reveal to Alice the real state  $s$ , thus confirming her belief that Heads is up. So in the model  $S$  from Example 1, we have  $s \models \Box_a H$ . In contrast, in Example 4, Alice’s belief (that the coin lies Heads up), though true, is *not safe*. There is some piece of correct information which, if learned by Alice, would make her change this belief: we can represent this piece of correct information as the doxastic proposition  $H \rightarrow K_b H$ . It is easy to see that the actual state  $s'_1$  (in the left-upper corner) of the model  $S'$  satisfies the proposition  $B_a^{H \rightarrow K_b H} T$ . So, if given this information, Alice would come to wrongly believe that the coin is Tails up! This is an example of a *dangerous truth*: a true information whose learning can lead to wrong beliefs.

**Degree of Safety of a Belief.** In dcpm’s, we can use the probabilistic information to refine our analysis of safe belief, by defining “degrees of safety” of a belief (similarly to the probabilistic degrees of belief). For any number  $x \in [0, 1]$ , we say that  $a$ ’s belief in  $Q$  has a *degree of safety of (at least)  $x$  at state  $s$* , and write  $s \in \Box_a^x Q$ , if  $a$ ’s degree of conditional belief in  $Q$  given  $P$  is at least  $x$  for all true conditions  $P$ . If we define the *degree of safety of  $a$ ’s belief in  $Q$  at state  $s$*  by

$$d_a^s(Q) := \min_{s \in P \subseteq s(a)} \mu_a(Q|P),$$

then  $\Box_a^x Q = \{s : d_a^s(Q) \geq x\}$ . “Safe belief”, as defined above, is the same as belief with degree of safety=1. A belief is “*weakly safe*” if it has degree of safety  $> 0$ : such a belief might be lost due to truthful learning, but it is never reversed (into believing the opposite). Indeed, it is easy to see that  $a$ ’s belief in  $Q$  is weakly safe iff  $\neg B_a^P \neg Q$  for all true propositions  $P$ . *Strongly unsafe* beliefs are the ones which (even if true) are not weakly safe (but have a null degree of safety). It is easy to see that, in Example 4 above, Alice’s belief that the coin lies Heads up is strongly unsafe. In many situations, it is enough to have a *high enough degree of safety* to pre-empt the potential dangers of learning.


**Degree of Common Safe Belief.** The standard notion of common belief can be extended to define a concept of “common safe belief of degree  $x$ ” (for  $0 \leq x \leq 1$ ):

$$C \Box^x P = \bigwedge_{a_1, a_2, \dots, a_n} \Box_{a_1}^x \Box_{a_2}^x \dots \Box_{a_n}^x P$$

(where the infinite conjunction ranges over all sequences  $a_1, \dots, a_n$  of agents). As for safe belief, we define *common safe belief*  $C \Box P$  simple as “common safe belief of degree 1”  $C \Box^1 P$ . One could argue that the above notions should play an important role in games: for instance, Aumann’s theorem [4] about backwards induction still holds if

we weaken his condition of “common knowledge of rationality” to “common safe belief of rationality”. Aumann’s celebrated result says that, if a state  $s$  satisfies *common knowledge of rationality*  $CkR$ , then  $s \in I$ . It is well-known that the theorem does *not* hold in the weaker assumption of *common true belief in rationality*. The reason is that players’ beliefs about the other players’ rationality *may change* during the game: they may lose belief in others’ rationality if some “surprising” moves are made.

**Example: The Centipedes Game.** As an example of what can go wrong, consider the following Centipedes Game, as presented in Board [12]:



We denote the nodes by the sequences of moves leading to them, so e.g. the original node is  $\emptyset$  and (as shown in the graph) player 1 is to move. Board provides a good analysis of this in terms of conditional beliefs [12]. In a nutshell, there exist plausibility models for this game such that the state  $(L_1T_3, L_3)$  satisfies in the same time  $B_2T_1$ ,  $B_2^{L_1} \neg R_1$  and  $B_2^{L_1L_2} L_3$  and common true belief in rationality. Intuitively, this is because player 2’s original belief that player 1 will play  $T_1$  (belief which is fully consistent with common true belief in rationality) is challenged by the surprising move  $L_1$ . To respond to this challenge, 2 must revise his beliefs, and it is perfectly possible (and consistent with 2’s rationality) that after this revision, 2 starts to believe that 1 is irrational, and moreover that 1 will play  $L_3$  if the node  $L_1L_2$  is reached. Given this belief, 2’s best response at node  $L_1$  is to play  $L_2$ . This explains the failure of backwards induction solution, i.e. the fact that  $(L_1T_3, L_3) \notin I$ .

To warrant the backwards induction solution, we would need common belief in rationality at all nodes that are *actually reachable*. To make this concept robust, the only uniform assumption that we can make at the original node  $\emptyset$  seems to be *common safe belief of rationality*. Using degrees of safety, we can refine this somewhat, when a specific game is given:

**Theorem 2.** If common safe belief in rationality holds at the initial state of a game of perfect information, then the backwards induction solution is played. More generally, for any game of perfect information  $G$ , there exists a number  $x > 0$  such that: if at the initial state of game  $G$  the degree of common safe belief in rationality is strictly bigger than  $x$ , then the backwards induction solution is played.

The smallest such  $x$  is called the *safety level* of game  $G$ , denoted by  $S(G)$ , and is computable in terms of the game tree  $G$ . For instance, a game has safety level 1 iff common safe belief in rationality holds at the initial state. But the safety

level of the Centipedes Game above is  $2/3$ : it is enough for player 2’s initial belief in 1’s rationality to have a degree of safety of more than  $2/3$  (i.e. he would always be cautious enough to assign only a probability of less than  $1/3$  to other player’s irrationality, no matter what new information he learns). This would give player 2 an expected utility strictly less than  $4/3 + 5/3 = 3$  if he plays  $L_2$ , i.e. less than the expected utility for playing  $T_2$ , thus motivating him to play the backwards induction solution. (But of course, for all he knows, player 2 might still be *wrong*: player 1 might actually *be irrational*, in which case player 2 misses a good opportunity to make a lot of money!)

**Theorem 3.** The logic of knowledge and safe belief introduced in [10], and presented in Appendix B below, is *complete with respect to dcpm’s*.

## 4 Action Models and Conditional Probabilistic Update


We now improve on the work in [3, 32, 10, 9, 11] by introducing *action models*, of both the conditional-probabilistic and the plausibilistic type, in order to represent *uncertain forms of multi-agent learning, belief updating or communication actions*: A (discrete conditional-probabilistic, or finite plausibility) *action model* is just a (discrete conditional-probabilistic, or finite plausibility) frame  $\Sigma$ , together with a *precondition map*

$$pre : \Sigma \rightarrow Prop$$


associating to each element of  $\Sigma$  some doxastic proposition  $pre_\sigma$ . As in [10, 9], we call the elements of  $\Sigma$  (*basic*) *doxastic actions*, and we call  $pre_\sigma$  the *precondition of action*  $\sigma$ . The precondition defines the *domain of applicability* of  $\sigma$ : this action can be executed on a state  $s$  iff  $s$  satisfies its precondition. The basic actions  $\sigma \in \Sigma$  are taken to represent some *deterministic* actions of a particularly simple nature. As mentioned in the Introduction, we only deal here with pure “belief changes”, i.e. actions that do not change the “ontic” facts of the world (but only the agents’ beliefs). Intuitively, a doxastic action represents the specific “learning event” that “triggers” the belief revision, while the action’s preconditions encode the *information* carried by the event. The conditional probabilities  $\mu_a$  (or the plausibility pre-orderings  $\leq_a$ ) give the *agent’s (probabilistic, or conditional) beliefs about the current learning event, at the moment of its happening*. As in [10, 9], these should be interpreted as *beliefs about changes*, that *encode changes of beliefs*. In this sense, we use such “beliefs about actions” simply as a way to represent doxastic changes.

**Example 4’.** Let us revisit Example 4, and think of the *action* leading to it: Bob takes a peek at the coin, when Alice goes away for a moment. Alice *believes that nothing is really happening* in her absence, though obviously she *doesn’t know* this (that nothing is happening). In the


literature on dynamic-epistemic logic, this action is usually called a *private announcement to a subgroup*: the “insider” (Bob) learns which face is up, while the outsider Alice believes nothing is happening. In the action model  $\Sigma$ , the two nodes on top represent the actions of Bob taking a peek and seeing that **H**, or respectively **T**, holds, while the notion the node on the bottom is the action in which “nothing is happening” (whose precondition is any tautology **true**):



**Example 5: Fully Successful Lying.** Suppose now that, after Bob took the peek in Alice’s absence (i.e. after we arrived in the situation described in Example 4), Bob sneakily announces: “Look, I took a peek and saw the coin was lying *Tails up*”. For our purposes, we can formalize the content of this announcement as  $K_b \mathbf{T}$ , i.e. saying that “Bob knows the coin is lying *Tails up*”. This is a *public announcement*, but *not a truthful one* (though it does convey some new truthful information): it is a *lie!* We assume that it is in fact a *fully successful lie*: it is common knowledge that, even after he admitted having taken a peek, Alice still trusts Bob, so she believes him. This action is given by the *left node* in the following model  $\Sigma'$ :



**Example 5': Partially Successful Lying.** In contrast, if admitting he cheated erodes somewhat Alice’s trust in Bob, so that she only assigns probability  $0 < x < 1$  to him telling the truth, then the action model is:



**Probabilistic Update Product.** To compute the output state model from the original (input) state model and the model of the action, we need a binary ‘update’ operation  $\otimes$ , taking any state model  $\mathbf{S} = (S, \triangleleft_a, \|\bullet\|)_{a \in \mathcal{A}}$  and any action model  $\Sigma = (\Sigma, \triangleleft_a, pre)_{a \in \mathcal{A}}$  into an *updated state model*  $\mathbf{S} \otimes \Sigma$ , representing the way an action lying in  $\Sigma$  will act on an input-state lying in  $\mathbf{S}$ . We call this the *update product* of the two models: as in [7, 6, 8, 10, 9], we take  $S \otimes \Sigma := \{(s, \sigma) : s \in pre(\sigma)_{\mathbf{S}}\}$  as the *set of states of the new model*<sup>13</sup>. For simplicity, we denote by  $s\sigma$  the pair  $(s, \sigma)$  seen as an output-state in  $\mathbf{S} \otimes \Sigma$ . The *valuation* is given by the original model<sup>14</sup>:  $s\sigma \models p$  iff  $s \models p$ . As in

<sup>13</sup>The justification is that: (1) basic actions are deterministic, so we can identify their outputs with pairs  $(s, \sigma)$  of an input and an action; and (2) an action is executable only on inputs that satisfy its precondition.

<sup>14</sup>This is because we only consider “purely doxastic” actions, so the ontic “facts” are left unchanged.

[7, 6], the new epistemic uncertainty relation is the *product* of the two uncertainty relations<sup>15</sup>:

$$s\sigma \sim_a s'\sigma' \quad \text{iff} \quad \sigma \sim_a \sigma', s \sim_a s'.$$

Furthermore, the new conditional probability is given in terms of the priority degrees, by:

$$(s\sigma, t\tau) = \lim_{x \rightarrow (s,t)} \frac{x \cdot (\sigma, \tau)}{x \cdot (\sigma, \tau) + (1-x) \cdot (\tau, \sigma)}$$

where the limit is taken over  $x$ ’s such that the denominator is  $\neq 0$ . This last clause can be justified by thinking of what doxastic actions *do*: they are events that *change* the prior (probabilistic) beliefs, via fresh, independent evidence that may *possibly override the prior beliefs* (when they flatly contradict the new evidence). The prior probability is not to be necessarily kept unchanged, since it reflects past beliefs; while the probabilities in the action model represent the agent’s *current beliefs* about . Independence would suggest multiplication of probabilities (followed by conditionalization on the action’s precondition (denominator 0), the prior beliefs not firmly held, but are assumed to be prone to small errors: so we extend the definition by continuity to the case that the denominator is 0. As in the AGM theory, an agent keeps as much as possible of his prior (certain) beliefs, as long as they are not contradicted by the new (certain) beliefs: *prior certainty can only be overridden if it is contradicted by current certainty; in addition, current certainty always overrides prior uncertainty; while prior uncertainty is “weighted” using the current uncertainty*.<sup>16</sup>

Spelling out the effect of the last clause in detail, we obtain the following equivalent definition by cases:

If  $(\sigma, \tau) = 0$ , then

$$(s\sigma, t\tau) = 0;$$

if  $(\sigma, \tau) = 1, \sigma \neq \tau$ , then

$$(s\sigma, t\tau) = 1;$$

otherwise


$$(s\sigma, t\tau) = \frac{(s, t) \cdot (\sigma, \tau)}{(s, t) \cdot (\sigma, \tau) + (1 - (s, t)) \cdot (1 - (\sigma, \tau))}.$$

In particular, the first two cases above give us the *Action-Priority Update* from [11] (also called *anti-lexicographic product update* in [9]): in plain words, the plausibility relation on the Cartesian product  $\mathbf{S} \times \Sigma$  of two plausibility models is taken to be the anti-lexicographic preorder induced by the two plausibility preorders.

<sup>15</sup>This encodes the intuition that the initial uncertainty about actions is *independent* of the initial uncertainty about states.

<sup>16</sup>A different justification can be provided using a generalization of Jeffrey’s Rule to Popper probabilities.


**Examples 5, 5' revisited.** We can see the qualitative difference between “fully successful lying” (which may completely overturn the agent’s prior beliefs) and the only “partially successful lying” (which only “weights” these prior beliefs), by comparing the way the actions in Examples 5 and 5' update the model  $S'$  in Example 4: the first action yields



while the second action leaves model  $S'$  essentially unchanged! So, in this sense, a “partially successful” lie is in fact partially successful only when the hearer already had some doubts to start with (in which case it enhances the doubt). It cannot override the hearer’s prior certainty; but a “fully successful” lie always does!

**Other Examples of Update Products.** It is easy to see that the update product of the state model in Example 1 and the action model in Example 4' is indeed (as expected) the state model in Example 4.


**Public Announcements of “Hard Facts”= Conditionalization.** A truthful public announcement  $P!$  of some “hard fact”  $P$  establishes common knowledge that  $P$  was the case. The action model consists of only one node, labeled with  $P$ . Its effect on a state model  $S$  (via the update product) is to delete all the non- $P$  states, keep the indistinguishability relations between the surviving states and change the probabilities conditionalizing with  $P$ :  $\mu'_a(Q|R) := \mu_a(Q|R \cap P_S)$ . An example of this is publicly announcing, in the situation described in Example 4, that if the coin lies Heads up then Bob knows it. This corresponds to  $(H \rightarrow K_b H)!$ , and the updated state model after that is:



Note that, in this model, Alice came to the (wrong!) belief that  $T$  (i.e. the coin lies Tails up): as we saw, this is only possible since her previous true belief that  $H$  was not safe.


**Public Announcement of “Soft” Facts.** Suppose an announcement  $P!_x$  is made, in such a way that all the agents believe with probability  $x$  it is truthful, although they don't

know for sure that it is truthful.



Note that the effect of such a “soft” announcement is different from the previous “hard” announcement, even when  $x = 1$ . The case  $x = 1$  has in fact been considered by other authors, who proposed a notion of “soft update” for it. But it is easy to see that its effect matches what we get by updating any given state model  $S$  with the action  $P!_1$  using our notion of product update: The new state model  $S \otimes P!_1$  can be thought of as being obtained from  $S$  by keeping the same information cells, and keeping the same priority order between any two states  $s, t \in P$ , and similarly between states  $s, t \notin P$ , while in the same time giving to all  $P$ -states priority with respect to all non- $P$  states.

**Discovery of Deceit.** Suppose that, in fact, when Bob was secretly taking a peek, Alice was taping him (using a hidden camera), so that she was able to see Bob taking a peek. Suppose also that it is common knowledge that Bob doesn't know that he is being videotaped, and that he assigns a very small probability  $1 - x$  to this possibility; in other words, he believes with some high probability  $x$  that there is no hidden camera. The action model is:



**Interception of messages** If the above “deceit” (Bob secretly looking at the coin) is replaced by a secret communication (from Charles to Bob, telling him that the coin lies Heads up), then the action model for “discovery of deceit” by Alice (as above) can also be interpreted as secret interception (wiretapping) by Alice of the secret message.

**Theorem 4.** A qualitative logic of belief-changing actions, with the same syntax as the one in [9, 11], can be developed<sup>17</sup> and completely axiomatized with respect to dcpm's. In the absence of common knowledge or common (safe) belief operators, this dynamic logic has the same expressive power as the “static” logic of knowledge and safe belief in the previous section. The reduction laws, presented in Appendix C below, are the same as the ones in [9, 11].

Moreover, these reduction laws can be extended to the quantitative logic of degrees of safety  $\square_a^x$ . But this needs

<sup>17</sup>The semantics would assign to qualitatively described syntactic actions any conditional probabilistic action model inducing the same plausibility model.



a language that includes linear inequalities  $\sum_i d_a(\varphi_i) \geq x$  involving linear combinations of degrees of safety, similar to the linear combinations of degrees of belief/knowledge used in [29]. Details will follow in future work.

**Open Problems and Future Work** An extended version of this paper will appear in *Studies in Logic*, with College Publications, London. We plan to write a separate paper on the game-theoretic applications of safe belief and degrees of safety, including the proof of Theorem 2. Another follow up includes the work on the above-mentioned reduction laws for the dynamic logic of probabilistic modal operators  $\Box_a^x Q$  (and  $B_a^{P,x} Q$ ), expressing that the degree of safety of the belief in  $Q$  (or the degree of conditional belief in  $Q$  given  $P$ ). We also list here three important open problems: (1) Study and axiomatize the corresponding logics for *infinite* conditional probability models. (2) Axiomatize the “static” logic of *probabilistic* modal operators  $\Box_a^x Q$  (and  $B_a^{P,x} Q$ ). (3) Axiomatize the *logic of common safe belief*  $C\Box Q$  and its quantitative version  $C\Box^x Q$ .

## Appendix A.

### A Complete Proof System for $CDL$

The *syntax* of  $CDL$  (without common knowledge and common belief operators) is:

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \varphi \mid B_a^\varphi \varphi$$

while the *semantics* over dcpm’s is given by the obvious compositional clauses (using the operators  $B_a^P Q$  and  $KQ$  defined in the paper). Note that in this logic, the *knowledge modality* is a derived operator, defined as an abbreviation by putting

$$K_a \varphi := B_a^{\neg\varphi} \varphi.$$

A *doxastic proposition* is a map  $\mathbf{P}$  assigning to each dcpm  $\mathbf{S}$  some  $\mathbf{S}$ -proposition, i.e. a set of states  $\mathbf{P}_\mathbf{S} \subseteq S$ . The interpretation map for the logic  $CDL$  can thus be thought of as associating to each sentence  $\varphi$  of  $CDL$  a doxastic proposition  $\|\varphi\|$ . We denote by *Prop* the family of all doxastic propositions. In addition to the rules and axioms of propositional logic, the *proof system* of  $CDL$  includes the following:

#### Necessitation Rule:

From  $\vdash \varphi$  infer  $\vdash B_a^\psi \varphi$ .

**Normality:**  $\vdash B_a^\theta(\varphi \rightarrow \psi) \rightarrow (B_a^\theta \varphi \rightarrow B_a^\theta \psi)$

**Truthfulness of Knowledge:**  $\vdash K_a \varphi \rightarrow \varphi$

**Persistence of Knowledge:**  $\vdash K_a \varphi \rightarrow B_a^\theta \varphi$

**Full Introspection:**  $\vdash B_a^\theta \varphi \rightarrow K_a B_a^\theta \varphi$   
 $\vdash \neg B_a^\theta \varphi \rightarrow K_a \neg B_a^\theta \varphi$

#### Hypotheses are (hypothetically) accepted:

$\vdash B_a^\varphi \varphi$

#### Minimality of revision:

$\vdash \neg B_a^\varphi \neg\psi \rightarrow (B_a^\varphi \wedge \psi \leftrightarrow B_a^\varphi(\psi \rightarrow \theta))$

## Appendix B.

### The Logic of Knowledge and Safe Belief

The *syntax* of the logic  $K\Box$  of knowledge and safe belief is:

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box_a \varphi \mid K_a \varphi$$

while the *semantics* over dcpm’s is given by the obvious compositional clauses. Note that in this logic, (conditional) belief is a derived operator, defined as an abbreviation:

$$B_a^\varphi \psi := \tilde{K}_a \varphi \rightarrow \tilde{K}_a(\varphi \wedge \Box_A(\varphi \rightarrow \psi)),$$

where

$$\tilde{K}_a \varphi := \neg K_a \neg \varphi$$

is the Diamond modality for  $K$ . This logic is more expressive than  $CDL$ .

The following set of axioms is complete for dcpm’s (and plausibility models):

- $K$ -axiom for  $K_a$  and  $\Box_a$ ;
- $S5$ -axioms for  $K_a$ ;
- $S4$ -axioms for  $\Box_a$ ;
- $K_a P \rightarrow \Box_a P$ ;
- $K_a(P \vee \Box_a Q) \wedge K_a(Q \vee \Box_a P) \rightarrow K_a P \vee K_a Q$ .

## Appendix C.

### The dynamic logic of doxastic actions

**Dynamic Modalities** Given a doxastic action  $\sigma$  (living in some action model  $\Sigma$ ), we can define a corresponding dynamic modality, capturing the *weakest precondition* of  $\sigma$ : for every proposition  $\mathbf{P}$ , the proposition  $[\sigma]\mathbf{P}$  is given by

$$([\sigma]\mathbf{P})_\mathbf{S} := \{s \in S : (s, \sigma) \text{ (if defined)} \in \mathbf{P}_{\mathbf{S} \otimes \Sigma}\}$$

**Syntax of Dynamic Logic of Doxastic Actions.** This was briefly sketched in [10]: As in [6], we consider a doxastic *signature*, i.e. a *finite (fixed) plausibility frame*  $\Sigma$ , together with an *ordered list without repetitions*  $(\sigma_1, \dots, \sigma_n)$  of some of the elements of  $\Sigma$ . Each signature gives rise to a dynamic-doxastic *logic*  $L(\Sigma)$ , as in [6]: one defines by double recursion a set of *sentences*  $\varphi$  and a set of *program terms*  $\pi$ ; the *basic programs* are of the form  $\pi = \sigma \vec{\varphi} = \sigma \varphi_1 \dots \varphi_n$ , where  $\sigma \in \Sigma$  and  $\varphi_i$  are sentences in our logic; program terms are generated from basic programs using *non-deterministic sum (choice)*  $\pi \cup \pi'$  and *sequential composition*  $\pi; \pi'$ . Sentences are built using the operators of the logic  $K\Box$  above, and in addition a dynamic modality  $[\pi]\varphi$ , taking program terms and sentences into other sentences. As in [6], the plausibility preorders on the signature  $\Sigma$  induce in a natural way plausibility preorders on basic programs in  $CDL(\Sigma)$ : we put

$(\sigma\vec{\varphi})_a^{\Pi\vec{\varphi}} := \{\sigma'\vec{\varphi} : \sigma' \in \sigma_a^{\Pi}\}$ . The given listing can be used to assign syntactic preconditions for basic programs, by putting:  $pre(\sigma_i\vec{\varphi}) := \varphi_i$ , and  $pre(\sigma\vec{\varphi}) := \top$  (the trivially true sentence) if  $\sigma$  is not in the listing. Thus, the basic programs of the form  $\sigma\vec{\varphi}$  form a (finite) *syntactic plausibility model*<sup>18</sup>  $\Sigma\vec{\varphi}$ . Every given interpretation  $\|\bullet\| : L(\Sigma) \rightarrow Prop$  of sentences as doxastic propositions will convert this syntactic model into a “real” (semantic) plausibility model, called  $\Sigma\|\vec{\varphi}\|$ .

To give the *semantics*, choose any dcpf  $(\Sigma, \mu)$  whose priority frame is isomorphic to  $\Sigma$ . We define by induction two *interpretation maps*, one taking any sentence  $\varphi$  to a doxastic proposition  $\|\varphi\| \in Prop$ , the second taking any program term  $\alpha$  to a (possibly non-deterministic) doxastic “program”, i.e. a *set* of basic actions in some dcpf. The definition uses the obvious semantic clauses and is completely similar to the one in [6], so we skip the details here.

The *proof system* is obtained by adding to the axioms and rules of the logic  $K\Box$  the following Reduction Axioms:

$$\begin{aligned}
[\alpha]p &\leftrightarrow pre_\alpha \rightarrow p \\
[\alpha]\neg\varphi &\leftrightarrow pre_\alpha \rightarrow \neg[\alpha]\varphi \\
[\alpha](\varphi \wedge \psi) &\leftrightarrow pre_\alpha \rightarrow [\alpha]\varphi \wedge [\alpha]\psi \\
[\alpha]K_a\varphi &\leftrightarrow pre_\alpha \rightarrow \bigwedge_{\beta \sim_a \alpha} K_a[\beta]\varphi \\
[\alpha]\Box_a\varphi &\leftrightarrow pre_\alpha \rightarrow \bigwedge_{\alpha \triangleleft_a \beta} K_a[\beta]\varphi \wedge \bigwedge_{\alpha \cong_a \gamma} \Box_a[\gamma]\varphi \\
[\pi \cup \pi']\varphi &\leftrightarrow [\pi]\varphi \wedge [\pi']\varphi \\
[\pi; \pi']\varphi &\leftrightarrow [\pi][\pi']\varphi
\end{aligned}$$

where  $p$  is any atomic sentence,  $\pi, \pi'$  are programs and  $\alpha$  is an *action*, i.e. a *basic program* in  $L(\Sigma)$ ,  $\sim_a$  is epistemic indistinguishability between actions,  $\triangleleft_a$  is *strict plausibility* order on actions, while  $\cong_a$  is *equi-plausibility* of (indistinguishable) actions:  $\alpha \cong_a \beta$  iff  $\alpha \triangleleft_a \beta$  and  $\beta \triangleleft_a \alpha$ .

## References

- [1] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [2] H. Arlo-Costa and R. Parikh. Conditional probability and defeasible inference. *Journal of Philosophical Logic*, 34:97–119, 2005.
- [3] G. Aucher. A combined system for update logic and belief revision. Master’s thesis, ILLC, University of Amsterdam, Amsterdam, the Netherlands, 2003.
- [4] R.J. Aumann. Backwards induction and common knowledge of rationality. *Games and Economic Behavior*, 8:6–19, 1995.
- [5] A. Baltag. A logic for suspicious players: epistemic actions and belief updates in games. *Bulletin Of Economic Research*, 54(1):1–46, 2002.
- [6] A. Baltag and L.S. Moss. Logics for epistemic programs. *Synthese*, 139:165–224, 2004. Knowledge, Rationality & Action 1–60.
- [7] A. Baltag, L.S. Moss, and S. Solecki. The logic of common knowledge, public announcements, and private suspicions. In I. Gilboa, editor, *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98)*, pages 43–56, 1998.
- [8] A. Baltag and S. Smets. Conditional doxastic models: a qualitative approach to dynamic belief revision. *Electronic Notes in Theoretical Computer Science*, 165:5–21, 2006.
- [9] A. Baltag and S. Smets. Dynamic belief revision over multi-agent plausibility models. In W. van der Hoek and M. Wooldridge, editors, *Proceedings of LOFT’06*, pages 11–24. Liverpool, 2006.
- [10] A. Baltag and S. Smets. The logic of conditional doxastic actions: a theory of dynamic multi-agent belief revision. In *Proceedings of ESSLLI Workshop on Rationality and Knowledge*. 2006.
- [11] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In W. van der Hoek G. Bonanno and M. Wooldridge, editors, *Texts in Logic and Games*. Amsterdam University Press, 2007. To appear.
- [12] O. Board. Dynamic interactive epistemology. *Games and Economic Behaviour*, 49:49–80, 2004.
- [13] C. Boutilier. On the revision of probabilistic belief states. *Notre Dame Journal of Formal Logic*, 36(1):158–183, 1995.
- [14] B. C. Van Fraassen. Representational of conditional probabilities. *Journal of Philosophical Logic*, 5:417–430, 1976.
- [15] B. C. Van Fraassen. Fine-grained opinion, probability, and the logic of full belief. *Journal of Philosophical Logic*, 24:349–377, 1995.
- [16] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge, MA, 1988.
- [17] J. Gerbrandy. Dynamic epistemic logic. *Logic, Language and Information*, 2, 1999. CSLI Publications, Stanford University.
- [18] J.Y. Halpern. Lexicographic probability, conditional probability, and nonstandard probability. In *Proceedings of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge (TARK 8)*, pages 17–30, 2001.
- [19] J.Y. Halpern. *Reasoning about Uncertainty*. MIT Press, Cambridge MA, 2003.
- [20] H. Katsuno and A. Mendelzon. On the difference between updating a knowledge base and revising it. In *Cambridge Tracts in Theoretical Computer Science*, pages 183–203, 1992.
- [21] B. P. Kooi. Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information*, 12:381–408, 2003.
- [22] G. Pappas and M. Swain, editors. *Essays on Knowledge and Justification*. Cornell Univ. Press, Ithaca, NY, 1978.
- [23] K.R. Popper. *The Logic of Scientific Discovery (revised Edition)*. London Hutchison, 1968. First Ed. 1934.
- [24] A. Renyi. On a new axiomatic theory of probability. *Acta Mathematica Academiae Scientiarum Hungaricae*, 6:285–335, 1955.
- [25] A. Renyi. Sur les espaces simples des probabilités conditionnelles. *Annales de L’institut Henri Poincaré: section B: calcul des probabilités et statistique*, B1:3–21, 1964.
- [26] R. Stalnaker. Knowledge, belief and counterfactual reasoning in games. *Economics and Philosophy*, 12:133–163, 1996.
- [27] J. van Benthem. Conditional probability meets update logic. *Journal of Philosophical Logic*, 12(4):409–421, 2003.
- [28] J. van Benthem. Dynamic logic of belief revision. Technical Report, ILLC, DARE electronic archive, University of Amsterdam, 2006.
- [29] J.F.A.K. van Benthem, J. Gerbrandy, and B.P. Kooi. Dynamic update with probabilities. In W. van der Hoek and M. Wooldridge, editors, *Proceedings of LOFT’06*. Liverpool, 2006.
- [30] J.F.A.K. van Benthem and F. Liu. Dynamic logic of preference upgrade. Technical report, 2004. ILLC Research Report PP-2005-29.
- [31] J.F.A.K. van Benthem, J. van Eijck, and B.P. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.
- [32] H.P. van Ditmarsch. Prolegomena to dynamic logic for belief revision. *Synthese (Knowledge, Rationality & Action)*, 147:229–275, 2005.

<sup>18</sup>A *syntactic plausibility model* is just a plausibility frame endowed with a *syntactic precondition map*, associating sentences to basic actions. For justification and examples, in the context of *epistemic action models*, see [6].