Varieties of Belief and Probability

Jan van Eijck [1]

CWI and ILLC
Amsterdam

Abstract

For reasoning about uncertain situations, we have probability theory, and we have logics of knowledge and belief. How does elementary probability theory relate to epistemic logic and the logic of belief? The paper focuses on the notion of betting belief, and interprets a language for knowledge and belief in two kinds of models: epistemic neighbourhood models and epistemic probability models. It is shown that the first class of models is more general in the sense that every probability model gives rise to a neighbourhood model, but not vice versa. The basic calculus of knowledge and betting belief is incomplete for probability models. These formal results were obtained in Van Eijck and Renne [9].

Keywords: Belief, betting, chance, foundations of subjective probability, Bayesian conditioning, neighbourhood models.

1 Introduction

Elementary probability theory, in the subjective or Bayesian style, is fascinating for cognitive scientists, for there is a marked contrast between fast error-prone assessment of chance and the slow but more accurate calculation of subjective probabilities using conditioning. Interest is added by the fact that belief about chance is an important basis of rational decision making and intelligent interaction. I know from our collaboration in the Games, Actions, and Social Software Project at NIAS that resulted in [11] and [12], that this is the stuff that Rineke loves.

Probability theorists like to view the difference between logic and probability as a difference in subject matter. Logic is the topic of reasoning about certainty, while probability theory teaches us how to reason about uncertainty. Guess which discipline has the most relevance to everyday life? Still, the probability theorists are right: epistemic or Bayesian probability can be viewed as an extension of propositional logic with hypotheses, i.e., basic propositions whose truth or falsity is uncertain. But logic has something to say, too, about reasoning under uncertainty: we have epistemic logic, doxastic logic, default

[1] Email: `jve@cwi.nl`

logic, logic of conditionals, and so on. So it is natural to ask how the perspectives of logic and probability theory on knowledge and belief are related. Frank Ramsey [27] considered the theory of probability as a branch of logic where arguments can be inconclusive. I wholeheartedly agree. In this paper I will argue that there is room for logics with a more general interpretation than probability measures. As an example of those, I explore neighbourhood models for a language of knowledge and belief as willingness to bet, and compare them with probabilistic models for the same language. The paper is light on formal definitions and proofs. For these, the reader is referred to Van Eijck and Renne [9].

The paper starts with some remarks, in Section 2, on the foundations of probability theory, as a comment on the views of Christiaan Huygens on probability. This is connected to the foundations of subjective probability on rational betting behaviour proposed by Ramsey, de Finetti and Savage, and, in Section 3, to the key role of probability in decision theory, which we owe to Von Neumann and Morgenstern. Section 4 introduces the notion of betting belief and compares this to some other notions of belief. In Section 5 I show that betting belief allows for a crisp analysis of the lottery puzzle, at the price of sacrificing closure of belief under conjunction. Section 6 presents a complete calculus for epistemic models with belief neighbourhoods, and Section 7 proves an incompleteness result for the calculus of betting belief with respect to probabilistic models. This shows that the logic of betting belief describes a more general kind of situation than is covered by probability models. Section 8 concludes.

2 Christiaan Huygens on the foundations of probability

Probability theory was invented by Pierre de Fermat and Blaise Pascal around 1650. The Dutch mathematician, astronomer, physicist and inventor Christiaan Huygens (1629–1695) picked up the new ideas during a visit to Paris in 1655. A digest of these was published, in Dutch, as an appendix to a textbook by a former mathematics teacher of Huygens, Frans van Schooten. This was the first treatise on probability that ever appeared, in Latin in 1657, and in Dutch in 1660. Its importance is in the game-theoretic foundation that Huygens proposes for probability, to support the technical results of Fermat and Pascal.

Huygens starts his essay on how to calculate what non-finished hazard games are worth and how to calculate winning chances in such games as follows:

> “Ick neeme tot beyder fondament, dat in het speelen de kansse, die yemant ergens toe heeft, even soo veel weerdt is als het geen, het welck hebbende hy weder tot deselfde kansse kan geraecken met rechtmatigh spel, dat is, daer in niemandt verlies geboden werdt. By exempel. So yemandt sonder mijn weeten in d’eene handt 3 schellingen verbergt, en in d’ander 7 schellingen, ende my te kiesen geeft welck van beyde ick begeere te hebben, ick segge dit my even soo veel weerdt te zijn, als of ick 5 schellingen seecker hadde.

> Om dat, als ik 5 schellingen hebbe, ick wederom daer toe kan geraecken, dat ick gelijcke kans sal hebben, om 3 of 7 schellingen te krijgen, en dat met rechtmatigh spel: gelijck hier naer sal betoont werden." [18]

Translation:

> "I take as the foundation of both [calculating what non-finished games are worth, and calculating winning chances] that in playing the chance that someone has in some matter, is worth just as much as the amount that, if he possesses it, will give him the same chances in a fair game, that is a game where no loss is offered to anyone. For instance. Suppose someone without my knowing hides in one hand 3 shillings, and in the other 7 shillings, and he offers me the choice between the two hands. Then I would say that this offer is worth the same as having 5 shillings for sure. Because, if I have 5 shillings, I can wager them in such manner that I have equal chances of getting 3 or 7 shillings, and that in a fair game, as will be explained hereafter."

Huygens explains this transformation to a symmetric game by applying it to his example:

> "Indien ick gelijcke kans heb om 3 te hebben of 7, soo is door dit Voorstel mijn kansse 5 weerdt; ende het is seecker dat ick 5 hebbende weder tot de selfde kansse kan geraecken. Want speelende om de selve tegen een ander die daer 5 tegen set, met beding dat de geene die wint den anderen 3 sal geven; soo is dit rechtmaetig spel, ende het blijckt dat ick gelijcke kans hebbe om 3 te hebben, te weeten, als ick verlies, of 7 indien ick win; want alsdan treck ick 10, daer van ick hem 3 geef."

Translation:

> "If I have equal chances to have 3 or 7, then by my Proposal this chance is worth 5; and it is sure that if I have 5, I will get to the same chance. Because putting 5 at stake against someone who stakes 5 against it, with condition that the one who wins will give the other 3, one has a fair game, and it becomes clear that I have equal chance of getting 3, namely, if I lose, or 7 if I win; because if I win I draw 10, of which I give 3 to him."

Thus, Huygens starts out from the expectation of a single individual in a lottery-like situation. He gives a reconstruction of this in terms of an n-person game, where n is the number of proposed chances, with equal stakes, and symmetric roles. Huygens argues that the value of the stakes equals the expectation. If a stake of value x buys me a ticket for a symmetric game with equal stakes that has the same outcomes as the lottery-like situation that we started out with, then it must be that the game and the lottery are worth the same. The Dutch mathematician Hans Freudenthal, in his review of Huygens' theory of probability, remarks that "Equal Chance" is validly defined as free choice for the player in a symmetric situation [15].

This is remarkably close to the famous Dutch book argument as a foundation of probability, proposed much later by Ramsey [27], de Finetti [14], and

Savage [29]. A Dutch book is a collection of bets (so it is not a book, and why it is called Dutch is unclear) that together represent either a sure win or a sure loss for the person who makes the bets, no matter how the situation turns out.

Take the case of equal chances of getting a and b again. Suppose this is offered to you as a symmetric game, at a price x that is different from $\frac{a+b}{2}$. Let G be the game where you get a if you win and b if you lose. Let G' be the game where you get b if you win and a if you lose. Then the only difference between G and G' is that the roles of the two players are reversed. So we may assume that you can enter into both games for the same price x. Now if $x < \frac{a+b}{2}$, what you should do is invest x in G and x in G', and play the games simultaneously. This costs you $2x$, and it yields $a + b$, so this is a Dutch book in your favour. If $x > \frac{a+b}{2}$, and you are willing to enter G and G' for the price x, then your investment of $2x$ will get you only $a + b$, so you are losing no matter what. There is a Dutch book against you.

Let us be a bit more precise about how Huygens would turn an individual choice situation with m possibilities $s_1, \ldots, s_m$, with revenues given by $L_1, \ldots, L_m$, into a stake distribution game G for m players. The stake x would be the same for every player. The game would match the players with the possibilities. The utility function would be given by: if player i draws s_j then i gets L_j. Obviously, the expectation for each player in this game is $\frac{\sum_{i=1}^{m} L_i}{m}$, so that should be the value of an individual stake. Also, the game is obviously symmetric, for all players have equal chances of getting each of the "prizes" $L_1, \ldots, L_m$.

Now replace the revenues by probabilities. Instead of $L_1, \ldots, L_m$ we have $p_1, \ldots, p_m$ with $\sum_{i=1}^{m} p_i = 1$. Nothing changes. The expectation in the game is $\frac{1}{m}$, so this should be the value of an individual stake. Anyone who can get a stake in the game for less that $\frac{1}{m}$ can set up a Dutch book, and anyone who is willing to enter the game for more than $\frac{1}{m}$ faces a Dutch book against him.

3 Belief and decision making

The following is a model for decision making under uncertainty that is widely used. An agent faces a choice between a finite number of possible courses of action, say $a_1, \ldots, a_n$. The agent is uncertain about the state of the world. Say she considers states $s_1, \ldots, s_m$ possible. Now suppose there is a table of consequences c, with $c(s_i, a_j)$ giving the consequences of performing action a_j in state s_i. How can the agent choose between the available actions in a rational way?

In the first place we should model the preferences of the agent. Let us suppose there is a preference ordering R on the consequences, with cRc' expressing that either the agent is indifferent between c and c', or the agent strictly prefers c to c'. Assume R is transitive and reflexive. Then define cPc' as $cRc' \wedge \neg c'Rc$, so that cPc' expresses that the agent strictly prefers c to c'. The relation P is transitive and irreflexive.

A utility function $u : C \to \mathbb{R}$ is said to *represent* R if u satisfies $u(c) \geq u(c')$ iff cRc'.

Von Neumann and Morgenstern [25] showed how to turn this into a tool for decision making if one adds a probability measure P on the state set. So assume $P(s_i) \geq 0$ and $\sum_{i=1}^{m} P(s_i) = 1$. Then a utility function u on the consequences induces a utility function U on the actions, by means of

$$U(a_j) = \sum_{i=1}^{m} P(s_i)u(s_i, a_j).$$

Now it is clear how a rational agent who disposes of (i) a utility function u representing her preferences and (ii) a probability measure on what she thinks is possible decides on what to do. Such an agent will perform the action a_j that maximizes $U(a_j)$.

This is the reason why expositions of probability theory often make strong claims about the applicability of their subject. Blitzstein and Hwang [7] list a number of possible applications of probability, and they close off with the application to Life in general:

> "Life is uncertain, and probability is the logic of uncertainty. While it isn't practical to carry out a formal probability calculation for every decision made in life, thinking hard about probability can help us avert some common fallacies, shed light on coincidences, and make better predictions."

This cheerful attitude to decision making engenders a particularly straightforward view of belief. I believe in φ if the odds in favour of φ are larger than $1:1$. Odds in favour of φ are calculated by means of $\frac{P(\varphi)}{P(\neg\varphi)}$. So I believe in φ if the subjective probability I assign to the truth of φ is larger than the subjective probability I assign to the truth of $\neg\varphi$. This is in fact the straightforward view that you should only believe propositions which have a probability greater than one half. Call this notion of belief betting belief.

4 Betting belief

The notion of betting belief has a number of remarkable properties. It is not closed under conjunction: it does not follow from the facts that $P(\varphi) > P(\neg\varphi)$ and $P(\psi) > P(\neg\psi)$ that $P(\varphi \wedge \psi) > P(\neg\varphi \vee \neg\psi)$. For suppose p, q, r are three propositions that are mutually exclusive and have the same probability. Then $P(p \vee q) > P(\neg p \wedge \neg q)$ and $P(q \vee r) > P(\neg q \wedge \neg r)$. From the fact that p, q, r are mutually exclusive it follows that $(p \vee q) \wedge (q \vee r)$ is equivalent to q. On the other hand, $P((p \vee q) \wedge (q \vee r)) = P(q) < P(\neg q)$. The following model gives a picture of this situation. The propositions p, q, r are mutually exclusive and have the same probability $\frac{1}{3}$. It is left to the reader to check in the picture that $P(p \vee q) = \frac{2}{3}$, $P(\neg p \wedge \neg q) = \frac{1}{3}$, $P(q \vee r) = \frac{2}{3}$, $P(\neg q \wedge \neg r) = \frac{1}{3}$.

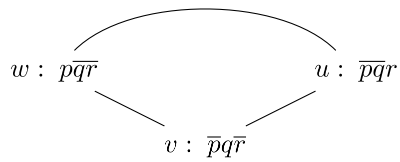

$$P(w) = P(v) = P(u) = \tfrac{1}{3}$$

This model represents probability by means of a weight function that gives each world the same weight. Note that the model also picture knowledge, which is represented by the epistemic accessibility relation.

The solid lines represent the epistemic accessibility relation of a single agent; they indicate that every world is accessible from any world. We will assume throughout this paper that knowledge accessibility is an equivalence; in other words, we are interpreting the knowledge operator K as an S5 operator. In the situation pictured above, the agent knows for instance that at least one of p, q, r is true. This is expressed by $K(p \vee q \vee r)$. The agent also knows that the propositions p, q, r are mutually exclusive. And so on.

Betting belief in φ and betting belief in $\varphi \to \psi$ does not entail betting belief in ψ. This is illustrated by the following model.

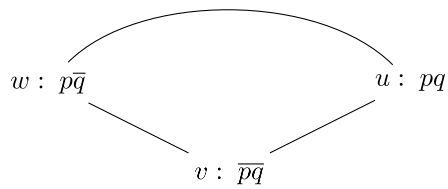

$$P(w) = P(v) = P(u) = \tfrac{1}{3}$$

Again, probability is represented by means of a weight function that gives each world the same weight. The probability of p (true in w and u) is $\frac{2}{3}$, the probability of $p \to q$ (true in v and u) is $\frac{2}{3}$, but the probability of q (true in u) is $\frac{1}{3}$. Thus, betting belief in p and $p \to q$ is justified, but betting belief in q is not.

On the other hand, betting belief in $p \wedge q$ implies betting belief in p and in q, for if the probability of $p \wedge q$ is greater than one half, then the same must hold for the probabilities of p and of q.

It is well known that people untrained in probability theory have difficulty with the notion of betting belief. Recall examples like the following. You are from a population with a statistical chance of 1 in 100 of having disease D. The initial screening test for this has a false positive rate of 0.2 and a false negative rate of 0.1. You tested positive; call this test result T. Should you believe you have the disease, with 'believe' in the sense of betting belief?

You reason: "If I test positive then, given that the test is quite reliable, the probability that I have D is quite high." So you tend to believe that you have D. But now you recall a lesson from your probability class: "True positives are often dwarfed by false positives." You pick up pen and paper and calculate:

$$P(D|T) = \frac{P(T|D)P(D)}{P(T)} = \frac{P(T|D)P(D)}{P(T|D)P(D) + P(T|\neg D)P(\neg D)}.$$

The first step uses Bayes' rule, and the second step calculates $P(T)$ by means of the rule of total probability. Filling in $P(T|D) = 0.9, P(D) = 0.01, P(\neg D) = 0.99, P(T|\neg D) = 0.2$, you arrive at the conclusion that $P(D|T) = \frac{1}{23}$.

As a result of your calculation, you don't believe anymore you have D but you agree to further testing. This step from an initial guess that the probability of D is high to a careful calculation revealing that the probability of D is low should perhaps be viewed as a switch from thinking fast to thinking slow, in the sense of Kahneman [21].

In any case, the example shows that qualitative belief judgements can be utterly misleading. Such examples made probability theorists like Richard Jeffrey urge us to give up qualitative belief altogether in favour of quantitative belief based on probability calculations [19,20]. In the rest of the paper I will show that there is room for qualitative belief linked to probability but not derived from it, after all.

The notion of betting belief introduced above can also be dubbed Bayesian belief. It is natural to interpret the uncertainties that we face in everyday life as subjective probabilities, and recalculations of betting belief based on new information can be viewed as model restrictions. The announcement φ in a model $\mathcal{M}$ leads to a new model $\mathcal{M}|\varphi$ consisting of all worlds in the old model that satisfy φ.

Consider the disease example again. Here is an epistemic probability model for it. The worlds are all connected, so this is a so-called S5 model. A weight function L gives the information about the probabilities for the four possible combinations of $d, \overline{d}$ with $t, \overline{t}$.

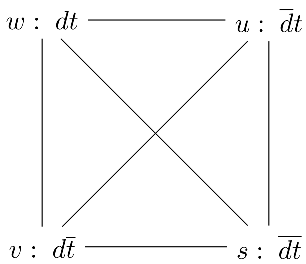

$$L(w) = 0.009, L(v) = 0.001, L(u) = 0.198, L(s) = 0.792$$

The weights (or probabilities, for the weight function is normalized) were computed by taking the prior probabilities for d, and multiplying with the appropriate error rates for the test. E.g., $L(u)$ is the product of $\frac{99}{100}$ (the prior probability of not having the disease) and $\frac{1}{5}$ (the false positive rate).

An update with the information t changes this model into the following restricted model, where the worlds where t is false have dropped out. This is the public announcement update from Jan Plaza [26].

$$w : \ dt \text{ ———— } u : \ \overline{d}t$$

$$L(w) = 0.009, L(u) = 0.198$$

Re-normalization of the weight function gives $L(w) = \frac{1}{23}, L(u) = \frac{22}{23}$. So after the information that the test was positive has been taken into account, the probability of d has changed from $\frac{1}{100}$ to $\frac{1}{23}$. The announcement update result agrees with the application of Bayes' rule. Bayesian conditionalisation (see [32]) and announcement update for epistemic probability models coincide. For further discussion and some qualification of this claim see [5].

Now let us coarsen the model, and replace the weight function by a neighbourhood function that tells us which propositions are believed in the betting sense. Starting out from epistemic models (Kripke models with equivalence relations of epistemic accessibility), we add a neighbourhood function for each epistemic agent. I will assume that within each i-cell, the neighbourhoods that get assigned to different worlds are the same; this encodes the fact that if an agent believes φ then she knows that she believes φ.

Truth definition for belief in φ, in terms of neighbourhoods, is:

$$M, w \models B\varphi \text{ iff for some } X \in N(w) \text{ for all } x \in X : M, x \models \varphi.$$

Here N is a function that assigns to each world w a set of neighbourhoods for w, where each neighbourhood X is a set of worlds. See [9] for a detailed comparison of neighbourhood models and epistemic probability models. Epistemic probability models are epistemic models with a weight function that assigns positive values to all worlds, and that satisfies the condition that the sum of the weights over each epistemic partition cell is bounded (but this condition is only relevant if the number of worlds in some partition cells is infinite).

Here is a neighbourhood version of the above epistemic weight model, with the neighbourhoods defined from the probabilities by means of: $X \in N(w)$ iff $X \subseteq [w]$ and $P(X) > P([w] - X)$, where $[w]$ is the epistemic equivalence class of w.

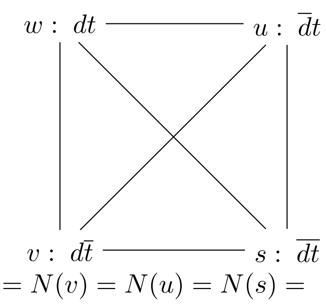

$$N(w) = N(v) = N(u) = N(s) = \{\{s\}, \{s,u\}, \{s,v\}, \{s,w\}, \{s,u,v\}, \{s,v,w\}, \{s,w,u\}, \{s,u,v,w\}\}.$$

To understand the neighbourhood function, observe first of all that since the epistemic accessibility relation is universal, the neighbourhoods are the same for every world. Next, note that that X is a neighbourhood iff $s \in X$. This is because the probability of world s in the original probability model is higher than the probability of $W - \{s\}$. It is convenient to use $\uparrow X$ for $\{Y \subseteq U \mid X \subseteq Y\}$ (the set of all supersets of X in domain U), where the domain U is understood from context. Then the neighbourhoods in the model are given by $N(w) = N(v) = N(u) = N(s) = \uparrow \{s\}$.

Now we can see that the neighbourhood function does not give enough information to calculate a new neighbourhood after information update. After information update with t, betting belief should favour world u over world w. But no reasonable update rule on neighbourhoods will give this result, for in the original model, the neighbourhood function is symmetric between w and u: we have for all neighbourhoods X that $w \in X$ iff $u \in X$.

This indicates that instead of a neighbourhood function we need something more expressive. One option here would be to introduce plausibility relations [2,3], and no doubt there are other options. The option we will explore here is modification of the neighbourhood function.

A *conditioned neighbourhood functional* is a functional $\mathfrak{N} : W \to \mathcal{P}(W) \to \mathcal{P}\mathcal{P}(W)$ that assigns to every w a function $\mathfrak{N}_w : \mathcal{P}(W) \to \mathcal{P}\mathcal{P}(W)$, where for each $X \subseteq W$, $\mathfrak{N}_w(X)$ is a set of neighbourhoods of w conditioned by X.

A neighbourhood functional for the disease model would assign to every world w and every $X \subseteq W$ a set of neighbourhoods given by

$$\mathfrak{N}_w(X) = \{Y \subseteq X \mid P(Y) > P(X - Y)\}.$$

For the disease model, we get the following neighbourhood functional (values indicated for all sets with size > 1):

$$\begin{array}{ll}
\{s,u,v,w\} & \mapsto \uparrow \{s\} \\
\{s,u,v\} & \mapsto \uparrow \{s\} \\
\{s,u,w\} & \mapsto \uparrow \{s\} \\
\{s,v,w\} & \mapsto \uparrow \{s\} \\
\{u,v,w\} & \mapsto \uparrow \{u\} \\
\{s,u\} & \mapsto \uparrow \{s\} \\
\{s,v\} & \mapsto \uparrow \{s\} \\
\{s,w\} & \mapsto \uparrow \{s\} \\
\{u,v\} & \mapsto \uparrow \{u\} \\
\{u,w\} & \mapsto \uparrow \{u\} \\
\{v,w\} & \mapsto \uparrow \{w\}
\end{array}$$

Truth definition for belief in φ, in terms of neighbourhood functionals is (assume $[w]$ gives the partition block of w for the epistemic relation):

$$M, w \models B\varphi \text{ iff for some } X \in \mathfrak{N}_w([w]) \text{ for all } x \in X : M, x \models \varphi.$$

A reasonable update rule for neighbourhood functionals could now be: restrict the functional to the new universe U.

Using this, we see that after update of the neighbourhood functional of the neighbourhood model with t, the agent still believes that $\neg d$, as she should.

One of the properties of betting belief is *strong commitment*. To see what that means, let us first look at the dual $\widehat{B}$ of B. $\widehat{B}\varphi$ is true iff $\neg B\neg\varphi$ is true iff it is not the case that the probability of $\neg\varphi$ is higher than $\frac{1}{2}$. This is the case iff the probability of $\neg\varphi \leq \frac{1}{2}$, iff the probability of φ is $\geq \frac{1}{2}$.

Now suppose $\widehat{B}\varphi$ is true. Then $P(\varphi) \geq \frac{1}{2}$. Suppose $\widehat{K}(\neg\varphi \wedge \psi)$ is also true. Then an accessible world where φ is false and ψ true exists. Let us look at the probability of $\varphi \vee \psi$. It must be strictly larger than $\frac{1}{2}$, for the world where φ is false and ψ true has positive weight. I have just shown the soundness of the following axiom of strong commitment (SC):

$$\widehat{B}\varphi \wedge \widehat{K}(\neg\varphi \wedge \psi) \to B(\varphi \vee \psi). \tag{SC}$$

Another axiom that we get immediately from the meaning of $\widehat{B}\varphi$ is (D) for determinacy:

$$B\varphi \to \widehat{B}\varphi. \tag{D}$$

What (D) says is that it follows from that fact that I am willing to bet on φ that I am not willing to bet on $\neg\varphi$.

If we replace the notion of betting belief by that of threshold belief, by interpreting belief in φ as $P(\varphi) > t$, for some specific t with $\frac{1}{2} \leq t < 1$ (this is also known as Lockean belief), then $\widehat{B}\varphi$ gets a different meaning. Under this notion of belief, $\widehat{B}\varphi$ is true iff it is not the case that $B\neg\varphi$ is true, iff it is not the case that $P(\neg\varphi) > t$, iff $P(\neg\varphi) \leq t$, iff $P(\varphi) \geq 1 - t$. Since $\frac{1}{2} \leq t$, this certainly holds. It follows that (D) also is sound for Lockean belief.

For threshold belief with $t > \frac{1}{2}$, (SC) fails, however. This is illustrated by the following counterexample.

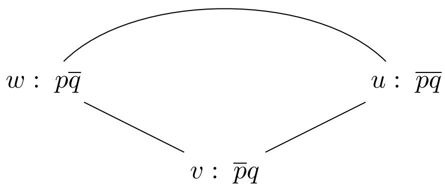

$$P(w) = 1 - t, P(v) = t - \tfrac{1}{2}, P(u) = \tfrac{1}{2}$$

Let $t > \frac{1}{2}$. Then $P(p) = P(w) = 1 - t$, so, as we have seen, $\widehat{B}p$ is true. Also, $\widehat{K}(\neg p \wedge q)$ is true, for there is an accessible world, v, where $\neg p \wedge q$ is true. The formula $p \vee q$ is true in worlds w and v, so $P(p \vee q) = P(w) + P(v) = 1 - t + t + \frac{1}{2} = \frac{1}{2} < t$, so $B(p \vee q)$ is false in the model.

Still another way to interpret (qualitative) belief is as follows: $S\varphi$ is true iff it holds for all consistent ψ that $P(\varphi|\psi) > P(\neg\varphi|\psi)$ (compare Leitgeb [24]). This uses $S\varphi$ for stable belief in φ. Stable belief can also be defined in terms of updates: $S\varphi$ is true in w iff it holds for all ψ that are true in w that $[!\psi]B\varphi$, where $[!\psi]B\varphi$ expresses that $B\varphi$ is true after updating the model with ψ, and where $B\varphi$ is interpreted as betting belief.

Since neighbourhood models are not expressive enough to model betting belief update, neighbourhood models cannot provide a reasonable truth definition for $S\varphi$. But if we switch to conditioned neighbourhood models, we have a means to interpret stable belief, as follows.

$$\begin{array}{lll} M, w \models S\varphi & \text{iff} & \text{for all } X \subseteq [w], X \neq \emptyset \\ & & \text{it holds for all } x \in X \text{ that } M \mid X, x \models B\varphi. \end{array}$$

Here $M \mid X$ is model M restricted to X, with the neighbourhood functional restricted accordingly, so $B\varphi$ is interpreted with respect to the updated neighbourhood functional. The clause for $S\varphi$ expresses that (stable) belief in φ is belief that continues to hold, no matter how we restrict the model.

In fact, Leitgeb's notion is a special case of this, for Leitgeb's theory is phrased in terms of standard Kripke models instead of neighbourhood models, and standard Kripke models can be viewed as constrained neighbourhood models.

Strong belief in φ, yet another notion of qualitive belief, is a bit harder to link to probability. It is defined for plausibility models, e.g., locally connected preorders. A preorder is a reflexive and transitive relation. A relation R is weakly connected (terminology of Robert Goldblatt [16]) if the following holds:

$$\forall x, y, z((xRy \wedge xRz) \rightarrow (yRz \vee y = z \vee zRy)).$$

A relation R is locally connected if both R and $R\check{}$ (the converse of R) are weakly connected. A most plausible possible world is a world that is maximal in the R ordering. An agent strongly believes in φ if φ is true in all most plausible accessible worlds. This yields a KD45 notion of belief (reflexive, euclidean, and serial). See Baltag & Smets [2,3].

Finally, it is possible to interpret qualitative belief as subjective certainty. An agent i believes in φ without any doubt if $P_i(\varphi) = 1$. This is used in epistemic game theory (Aumann [1]), and can easily be expressed in epistemic models, for this notion coincides with knowledge. If one drops the requirement that weight functions assign strictly positive values to all worlds then certainty and knowledge no longer coincide.

5 The lottery puzzle

One of attractions of betting belief lies in the light it sheds on the lottery puzzle. If Alice believes of each of the tickets 000001 through 111111 that they are not winning, then this situation is described by the following formula:

$$\bigwedge_{t=000001}^{111111} B_a \neg t.$$

If her beliefs are closed under conjunction, then this follows:

$$B_a \bigwedge_{t=000001}^{111111} \neg t.$$

But actually, she believes, of course, that one of the tickets is winning:

$$B_a \bigvee_{t=000001}^{111111} t.$$

This is a contradiction. Since the lottery puzzle involves three statements, there are three possible strategies to deal with it.

(i) Deny that Alice believes that her ticket is not winning.

(ii) Block the inference from $\bigwedge_{t=000001}^{111111} B_a \neg t$ to $B_a \bigwedge_{t=000001}^{111111} \neg t$.

(iii) Deny that Alice believes that there is a winning ticket.

A notion of belief for which it holds that Alice does not believe there is a winning ticket will hardly convince anyone, so let us forget about that way out. This leaves us with two options.

The advantage of (i) is that there is no need to sacrifice closure of belief under conjunction. A disadvantage is that one has to opt for a severe restriction of what counts as belief.

An advantage of (ii): no need to artificially restrict what counts as belief. And true, one has to sacrifice closure of belief under conjunction, but this is maybe not so bad after all. As I will see below, lots of nice logical properties remain.

Proponents of (i) are many philosophers, and they are easy to recognize: they call the lottery puzzle *the lottery paradox*. But maybe this is a bit harsh on the philosophers; after all, some have taken the trouble to develop notions of stable belief where some version of (i) can be saved. Proponents of (ii) are subjective Probabilists like Jeffrey [20], and decision theorists like Kyburg [23]. As we will see in the next section, one can side with them without giving up reasonable notions of qualitative belief.

6 Neighbourhood models and completeness

To drop the closure of belief under conjunction, we need an operator B_a that does *not* satisfy (Dist).

$$B_a(\varphi \to \psi) \to B_a\varphi \to B_a\psi \qquad \text{(Dist-B)}$$

This means: B_a is not a *normal* modal operator. See also [34]. Interpreting modal operators as accessibility relations between worlds brings the distribution axiom or K axiom in its wake. In order to drop it we have to switch to (epistemic) neighbourhood models. Here is a formal definition.

An **Epistemic Doxastic Neighbourhood Model** $\mathcal{M}$ for set of agents *Ag* and set of propositions *Prop* is a tuple

$$(W, R, V, N)$$

where

- W is a non-empty set of worlds.
- R is a function that assigns to every agent $a \in Ag$ an equivalence relation $\sim_a$ on W. We use $[w]_a$ for the $\sim_a$ class of w, i.e., for the set $\{v \in W \mid w \sim_a v\}$.
- V is a valuation function that assigns to every $w \in W$ a subset of *Prop*.
- N is a function that assigns to every agent $a \in Ag$ and world $w \in W$ a collection $N_a(w)$ of sets of worlds—each such set called a *neighbourhood* of w—subject to a set of *conditions*.

The core conditions are as follows:

(c) $\forall X \in N_a(w) : X \subseteq [w]_a$. This ensures that agent a does not believe any propositions $X \subseteq W$ that she knows to be false.

(f) $\emptyset \notin N_a(w)$. This ensures that no logical falsehood is believed.

(n) $[w]_a \in N_a(w)$. This ensures that what is known is also believed.

(a) $\forall v \in [w]_a : N_A(v) = N_A(w)$. This ensures that if X is believed, then it is known that X is believed.

By dropping some of these conditions one can further weaken (or: generalize, depending on perspective) the notion of belief. But the constraints that the conditions impose on belief are quite weak, so we will not do so here.

There are three further conditions that may be imposed to further strengthen the notion of belief.

(m) $\forall X \subseteq Y \subseteq [w]_A :$ if $X \in N_A(w)$, then $Y \in N_A(w)$. This says that belief is monotonic: if an agent believes X, then she believes all propositions $Y \supseteq X$ that follow from X. This may seem entirely reasonable, but in proposals where neighbourhoods are used to model conflicting and inconclusive evidence [6] it is dropped.

(d) If $X \in N_a(w)$ then $[w]_a - X \notin N_a(w)$. This corresponds to the axiom (D) that we discussed above. This condition says that if a believes a proposition X then a does not believe the negation of that proposition. As we have seen, this holds for betting belief and threshold belief, for a threshold above $\frac{1}{2}$. For thresholds below $\frac{1}{2}$, it fails, however.

(sc) $\forall X, Y \subseteq [w]_a$: if $[w]_a - X \notin N_a(w)$ and $X \subsetneq Y$, then $Y \in N_a(w)$. If the agent does not believe the complement $[w]_a - X$, then she must believe any strictly weaker Y implied by X. We saw above that this distinguishes betting belief from threshold beliefs for thresholds above $\frac{1}{2}$.

Epistemic doxastic neighbourhood models can interpret the language of epistemic doxastic logic (henceforth, KB language):

$$\varphi ::= \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid B_a\varphi.$$

The interpretation of $K_a\varphi$ uses the R relations; the interpretation of $B_a\varphi$ uses the neighbourhoods. Here is the neighbourhood version of the first example above:

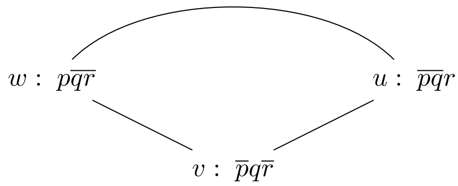

$$N(w) = N(v) = N(u) = \{\{w,v\},\{v,u\},\{w,u\},\{w,v,u\}\}$$

In all worlds, $K(p \vee q \vee r)$ is true. In all worlds $B\neg p$, $B\neg q$, $B\neg r$ are true. In all worlds $B(\neg p \wedge \neg q)$, $B(\neg p \wedge \neg r)$, $B(\neg q \wedge \neg r)$ are false. So the lottery puzzle is solved in neighbourhood models for belief by non-closure of belief under conjunction.

Here is a calculus for betting belief that relates belief to a standard S5 notion of knowledge.

AXIOMS

(Taut)	All instances of propositional tautologies
(Dist-K)	$K_a(\varphi \to \psi) \to K_a\varphi \to K_a\psi$
(T)	$K_a\varphi \to \varphi$
(PI-K)	$K_a\varphi \to K_aK_a\varphi$
(NI-K)	$\neg K_a\varphi \to K_a\neg K_a\varphi$
(F)	$\neg B_a\bot$.
(PI-KB)	$B_a\varphi \to K_aB_a\varphi$
(NI-KB)	$\neg B_a\varphi \to K_a\neg B_a\varphi$
(KB)	$K_a\varphi \to B_a\varphi$
(M)	$K_a(\varphi \to \psi) \to B_a\varphi \to B_a\psi$
(D)	$B_a\varphi \to \neg B_a\neg\varphi$.
(SC)	$\widehat{B}_a\varphi \wedge \widehat{K}_a(\neg\varphi \wedge \psi) \to B_a(\varphi \vee \psi)$

RULES

$$\frac{\varphi \to \psi \quad \varphi}{\psi}\ \text{(MP)} \qquad \frac{\varphi}{K_a\varphi}\ \text{(Nec-K)}$$

This calculus for betting belief is discussed in [9] and [4]. The fact that closes off this section is proved in [9].

Fact 6.1 *The calculus of betting belief is complete for epistemic doxastic neighbourhood models.*

7 Incompleteness for epistemic probability models

The step from neighbourhoods to probabilities is very small, but we will see in this section that the logic of neighbourhoods and the logic of probabilities are different.

Epistemic probability models are the result of replacing the neighbourhood function of an epistemic doxastic neighbourhood model by a weight function L. A weight function L assigns to every agent a a function $L_a : W \to \mathbb{Q}^+$ (the positive rationals), subject to the constraint that the sum of the L_a values over each epistemic partition cell of a is bounded. If $X \subseteq W$ then let $L_a(X)$ be shorthand for $\sum_{x \in X} L_a(x)$. Boundedness can then be expressed as follows: for each i and w: $L_a([w]_a) < \infty$.

To illustrate, here is an example from investment banking. Two bankers i, j consider buying stocks in three firms a, b, c that are involved in a takeover bid. There are three possible outcomes: a for "a wins", b for " b wins", and c for "c wins." i takes the winning chances to be $3 : 2 : 1$, j takes them to be $1 : 2 : 1$. In the following picture, the knowledge of i is represented by solid lines, that of j by dashed lines.

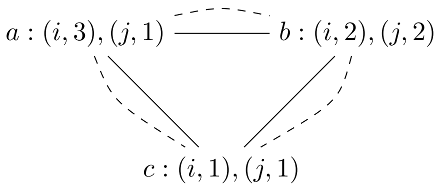

In all worlds, i assigns probability $\frac{1}{2}$ to a, $\frac{1}{3}$ to b and $\frac{1}{6}$ to c, while j assigns probability $\frac{1}{4}$ to a and to c, and probability $\frac{1}{2}$ to b.

We see that i is willing to bet $1 : 1$ on a, while j is willing to bet $3 : 1$ against a. It follows that in this model i and j have an opportunity to gamble, for, to put it in Bayesian jargon, they do not have a common prior.

Now consider the possibility that agent j has learnt something. Suppose that, as a result of this information, agent j (dashed lines) now considers c impossible.

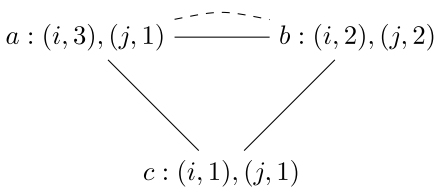

So we suppose that j has foreknowledge about what firm c will do.

The probabilities assigned by i remain as before. The probabilities assigned by j have changed, as follows. In worlds a and b, j assigns probability $\frac{1}{3}$ to a

and $\frac{2}{3}$ to b. In world c, j is sure of c.

We may suppose that this new model results from j being informed about the truth value of c, while i is aware that j received this information, but without i getting the information herself. So i is aware that j's subjective probabilities have changed, and it would be unwise for i to put her beliefs to the betting test. For although i cannot distinguish the three situations, she knows that j can distinguish the c situation from the other two. Willingness of j to bet against a at any odds can be interpreted by i as an indication that c is true, thus forging an intimate link between action and information update. I leave further analysis for another occasion.

Here is an example where two agents i (solid lines) and j (dashed lines) are uncertain about the toss of a coin. i holds it possible that the coin is fair f and that it is biased $\overline{f}$, with a bias $\frac{2}{3}$ for heads h. j can distinguish f from $\overline{f}$. The two agents share the same weight (so this is a single weight model, see [10]), and the weight values are indicated as numbers in the picture.

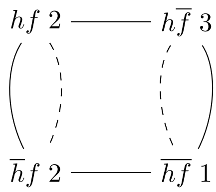

In world hf, i assigns probability $\frac{5}{8}$ to h and probability $\frac{1}{2}$ to f, and j assigns probability $\frac{1}{2}$ to h and probability 1 to f.

It is possible to normalize this model, but as a result of this each agent will have to get its own weight, for the weight functions are normalized within the epistemic accessibility cells.

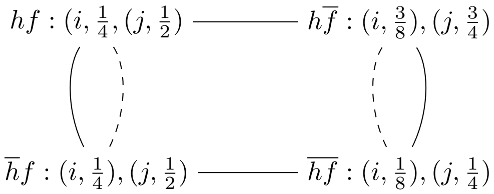

The rules for interpretation of the KB language in epistemic probability models are obvious:

$$\mathcal{M}, w \models K_a\varphi \text{ iff } \quad \text{for all } v \in [w]_a : \mathcal{M}, v \models \varphi.$$

$$\mathcal{M}, w \models B_a\varphi \text{ iff}$$

$$\sum\{L_a(v) \mid v \in [w]_a, \mathcal{M}, v \models \varphi\} > \sum\{L_a(v) \mid v \in [w]_a, \mathcal{M}, v \models \neg\varphi\}.$$

There is also an obvious way to reduce an epistemic probability model to a neighbourhood model, while preserving betting belief. Let $\mathcal{M} = (W, R, V, N)$

be a neighbourhood model and let L be a weigth function for $\mathcal{M}$. Then L *agrees with* $\mathcal{M}$ if it holds for all agents a and all $w \in W$ that

$$X \in N_a(w) \text{ iff } L_a(X) > L_a([w]_a - X).$$

The following theorem may come as a surprise, for it shows that, in a sense, the class of epistemic doxastic neighbourhood models is more general than that of probabilistic models. In other words: the principles of betting belief given in the calculus above do not force a probabilistic interpretation of the B operator.

Theorem 7.1 *There exists an epistemic doxastic neighbourhood model* $\mathcal{M}$ *that has no agreeing weight function.*

Proof. The proof of this uses an adaptation of an example from [33, pp. 344-345]. Let *Prop* $:= \{a, b, c, d, e, f, g\}$. Assume a single agent 0. Define:

$$\mathcal{X} := \{efg, abg, adf, bde, ace, cdg, bcf\}.$$

$$\mathcal{X}' := \{abcd, cdef, bceg, acfg, bdfg, abef, adeg\}.$$

Notation: xyz for $\{x, y, z\}$.

$$\mathcal{Y} := \{Y \mid \exists X \in \mathcal{X} : X \leq Y \leq W\}.$$

Let $\mathcal{M} := (W, R, V, N)$ be defined by $W :=$ *Prop*, $R_0 = W \times W$, $V(w) = \{w\}$, and for all $w \in W$, $N_0(w) = \mathcal{Y}$. Check that $\mathcal{X}' \cap \mathcal{Y} = \emptyset$. So $\mathcal{M}$ is a neighbourhood model.

Toward a contradiction, suppose there exists a weight function L that agrees with $\mathcal{M}$. Since each letter $p \in W$ occurs in exactly three of the seven members of $\mathcal{X}$, we have:

$$\sum_{X \in \mathcal{X}} L_0(X) = \sum_{p \in W} 3 \cdot L_0(\{p\}).$$

Since each letter $p \in W$ occurs in exactly four of the seven members of $\mathcal{X}'$, we have:

$$\sum_{X \in \mathcal{X}'} L_0(X) = \sum_{p \in W} 4 \cdot L_0(\{p\}).$$

On the other hand, from the fact that $L_0(X) > L_o(W - X)$ for all members X of $\mathcal{X}$ we get:

$$\sum_{X \in \mathcal{X}} L_0(X) > \sum_{X \in \mathcal{X}} L_0(W - X) = \sum_{X \in \mathcal{X}'} L_0(X).$$

Contradiction. So no such L_0 exists. □

I conjecture that this is the smallest counterexample, that is, I guess that all neighbourhood models up to size 6 have an agreeing weight function, but this needs to be checked.

Fact 7.2 *The calculus of epistemic-doxastic neighbourhood logic is sound for the class of epistemic probability models. Probabilistic beliefs are neighbourhoods.*

Theorem 7.1 shows that the KB calculus is incomplete for the class of epistemic probability models. In order to get a calculus that fits this class, we have to add an infinite series of axioms. The idea behind these axioms is from Scott [30]. What the axioms say, intuitively: If agent a knows the number of true φ_i is less than or equal to the number of true ψ_i, and if a believes φ_1, and the remaining φ_i are each consistent with her beliefs, then agent a believes one of the ψ_i.

It turns out that this is expressible in the KB language; see Segerberg [31]. Let $(\varphi_1, \dots, \varphi_m \mathbb{I}_a \psi_1, \dots, \psi_m)$ abbreviate the KB formula expressing that agent a knows that the number of true φ_i is less than or equal to the number of true ψ_i. Put another way, $(\varphi_i \mathbb{I}_a \psi_i)_{i=1}^m$ is true if and only if every one of a's epistemically accessible worlds satisfies at least as many ψ_i as φ_i. Using this, we can express the Scott axioms:

(Scott) $\quad [(\varphi_i \mathbb{I}_a \psi_i)_{i=1}^m \wedge B_a\varphi_1 \wedge \bigwedge_{i=2}^m \widehat{B}_a\varphi_i] \to \bigvee_{i=1}^m B_a\psi_i$

Theorem 7.3 *Adding the Scott axioms to the KB calculus yields a system that is sound and complete for epistemic probability models.*

For the proof of this I refer to [9]. To say a bit more about the connection between qualitative belief and quantitative belief we need a more expressive language for interpretation in epistemic probability models.

Let i range over *Ag*, p over *Prop*, and q over $\mathbb{Q}$. Then the language of epistemic probability logic is given by:

$$\varphi ::= \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid t_a \geq 0 \mid t_a = 0$$

$$t_a ::= q \mid q \cdot P_a\varphi \mid t_a + t_a \text{ where all indices } a \text{ are the same.}$$

This is expressive enough to compare subjective probabilities of the same agent. In particular, we can say things like $P_a(\varphi) > P_a(\psi)$. Truth for this language in epistemic probability models is defined as follows. Let $\mathcal{M} = (W, R, V, L)$ be an epistemic weight model and let $w \in W$.

$$\mathcal{M}, w \models \top \quad \text{always}$$

$$\mathcal{M}, w \models p \text{ iff } p \in V(w)$$

$$\mathcal{M}, w \models \neg\varphi \text{ iff it is not the case that } \mathcal{M}, w \models \varphi$$

$$\mathcal{M}, w \models \varphi_1 \wedge \varphi_2 \text{ iff } \mathcal{M}, w \models \varphi_1 \text{ and } \mathcal{M}, w \models \varphi_2$$

$$\mathcal{M}, w \models t_a \geq 0 \text{ iff } [\![t_a]\!]_w^{\mathcal{M}} \geq 0$$

$$\mathcal{M}, w \models t_a = 0 \text{ iff } [\![t_a]\!]_w^{\mathcal{M}} = 0.$$

$$[\![q]\!]_w^{\mathcal{M}} := q$$

$$[\![q \cdot P_a\varphi]\!]_w^{\mathcal{M}} := q \times P_{a,w}^{\mathcal{M}}(\varphi)$$

$$[\![t_a + t'_a]\!]_w^{\mathcal{M}} := [\![t_a]\!]_w^{\mathcal{M}} + [\![t'_a]\!]_w^{\mathcal{M}}$$

$$P_{a,w}^{\mathcal{M}}(\varphi) = \frac{L_a(\{u \in [w]_a \mid \mathcal{M}, u \models \varphi\})}{L_a([w]_a)}.$$

Fact 7.4 *A sound and complete calculus for the language of epistemic probability logic, interpreted in epistemic probability models, is given in [10].*

See also [13] and [22], where calculi for different epistemic probability model classes are given.

Notice that every epistemic probability model has an associated neighbourhood model. For if $\mathcal{M} = (W, R, V, L)$ is an epistemic probability model, then let $\mathcal{M}^\bullet$ be the tuple (W, R, V, N) given by replacing the weight function by a function N, where N is defined as follows, for $a \in Ag$, $w \in W$.

$$N_a(w) = \{X \subseteq [w]_a \mid L_a(X) > L_a([w]_a - X)\}.$$

Fact 7.5 *For any epistemic weight model $\mathcal{M}$ it holds that $\mathcal{M}^\bullet$ is a neighbourhood model.*

Now let us translate knowledge and belief into probability statements, by interpreting knowledge as certainty and belief as betting belief.

If φ is a KB formula, then $\varphi^\bullet$ is the formula of the language of epistemic probability logic given by the following instructions:

$$\begin{aligned}
\top^\bullet &= \top \\
p^\bullet &= p \\
(\neg\varphi)^\bullet &= \neg\varphi^\bullet \\
(\varphi_1 \wedge \varphi_2)^\bullet &= \varphi_1^\bullet \wedge \varphi_2^\bullet \\
(K_a\varphi)^\bullet &= P_a(\varphi^\bullet) = 1 \\
(B_a\varphi)^\bullet &= P_a(\varphi^\bullet) > P_a(\neg\varphi^\bullet).
\end{aligned}$$

Theorem 7.6 *For all KB formulas φ, for all epistemic probability models $\mathcal{M}$, for all worlds w of $\mathcal{M}$:*

$$\mathcal{M}^\bullet, w \models \varphi \textit{ iff } \mathcal{M}, w \models \varphi^\bullet.$$

Proof. Induction on formula structure. □

Theorem 7.7 *Let $\vdash$ denote derivability in the neighbourhood calculus for KB. Let $\vdash'$ denote derivability in the calculus of EPL. Then $\vdash \varphi$ implies $\vdash' \varphi^\bullet$.*

Proof. Induction on proof structure. □

8 Some Loose Ends

Are there applications where neighbourhoods without agreeing weight functions are natural? Is there a natural interpretation for the incompleteness example for $\{a, b, c, d, e, f, g\}$? Is the counterexample against completeness of the KB calculus for probability models the smallest counterexample?

Representation of probability information by means of weight functions was designed with implementation of model checking in mind. Just extend epistemic model checkers for S5 logics with a weight table for each agent. Implementations of model checkers for these logics can be found in [8] and in

[28]. The implementations can deal with Monty Hall style puzzles, urn puzzles, Bayesian updating by drawing from urns or tossing (possibly biased) coins, and 'paradoxes' such as the puzzle of the three prisoners (see, e.g., [20]). Efficiency was not a goal, but these implementations can be made quite efficient with a little effort.

Further analysis of the connection between neighbourhood logics and probabilistic logics [9] is in order. This is also connected to work of Wes Holliday and Thomas Icard [17]. Holliday and Icard investigate a language with a primitive operation $\varphi \succsim_a \psi$, for "according to a, φ is at least as probable as ψ." This is a revival of Segerberg's modal logic for comparative probability [31]. Interestingly, the qualitative probability Kripke models defined by Segerberg (and adopted by Holliday and Icard) seem better suited for defining well-behaved model restriction operations than the neighbourhood models used in the present paper. But note that the models with conditional neighbourhood functionals remedy this. Therefore, an obvious next step in the investigation of the logic of knowledge and qualitative belief is the study of the class of epistemic doxastic models with conditional neighbourhood functionals, together with operations of knowledge and belief update.

Acknowledgement

I thank the participants of the Synthese Workshop on Qualitative and Quantitative Methods in Formal Epistemology (Amsterdam, November 2014) for stimulating feedback. Two anonymous reviewers provided helpful comments on an earlier version of this paper.

References

[1] Aumann, R., *Interactive epistemology I: Knowledge*, International Journal of Game Theory **28** (1999), pp. 263–300.

[2] Baltag, A. and S. Smets, *Conditional doxastic models: A qualitative approach to dynamic belief revision*, Electronic Notes in Theoretical Computer Science (ENTCS) **165** (2006), pp. 5–21.

[3] Baltag, A. and S. Smets, *A qualitative theory of dynamic interactive belief revision*, in: G. Bonanno, W. van der Hoek and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, Texts in Logic and Games, Amsterdam University Press, 2008 pp. 11–58.

[4] Baltag, A., J. van Benthem, J. van Eijck and S. Smets, *Reasoning about communication and action* (2014), book manuscript, ILLC.

[5] Benthem, J. v., *Conditional probability meets update logic*, Journal of Logic, Language and Information **12** (2003), pp. 409–421.

[6] Benthem, J. v. and E. Pacuit, *Dynamic logics of evidence-based beliefs*, Studia Logica **99** (2011), pp. 61–92.

[7] Blitzstein, J. K. and J. Hwang, "Introduction to Probability," CRC Press, 2014.

[8] Eijck, J. v., *Learning about probability* (2013), available from `http://homepages.cwi.nl/~jve/software/prodemo`.

[9] Eijck, J. v. and B. Renne, *Belief as willingness to bet*, E-print, arXiv.org (2014), arXiv:1412.5090v1 [cs.LO].

[10] Eijck, J. v. and F. Schwarzentruber, *Epistemic probability logic simplified*, in: R. Goré, B. Kooi and A. Kurucz, editors, *Advances in Modal Logic, Volume 10*, 2014, pp. 158–177.

[11] Eijck, J. v. and R. Verbrugge, editors, "Discourses on Social Software," Texts in Logic and Games **5**, Amsterdam University Press, Amsterdam, 2009.

[12] Eijck, J. v. and R. Verbrugge, editors, "Games, Actions, and Social Software," Texts in Logic and Games, LNAI **7010**, Springer Verlag, Berlin, 2012.

[13] Fagin, R. and J. Halpern, *Reasoning about knowledge and probability*, Journal of the ACM (1994), pp. 340–367.

[14] Finetti, B. d., *La prevision: ses lois logiques, se sources subjectives*, Annales de l'Institut Henri Poincaré **7** (1937), pp. 1–68, translated into English and reprinted in Kyburg and Smokler, Studies in Subjective Probability (Huntington, NY: Krieger; 1980).

[15] Freudenthal, H., *Huygens' foundations of probability*, Historia Mathematica **7** (1980), pp. 113–117.

[16] Goldblatt, R., "Logics of Time and Computation, Second Edition, Revised and Expanded," CSLI Lecture Notes **7**, CSLI, Stanford, 1992 (first edition 1987), distributed by University of Chicago Press.

[17] Holliday, W. H. and T. F. Icard, *Measure semantics and qualitative semantics for epistemic modals*, in: *Proceedings of SALT*, SALT 23, 2013, pp. 514—-534.

[18] Huygens, C., "Van Rekeningh in Spelen van Geluck," 1660, about Calculation in Hazard Games; Latin: De ratociniis in Ludo aleae.

[19] Jeffrey, R., "The Logic of Decision," University of Chicago Press, 1983, second edition.

[20] Jeffrey, R., "Subjective Probability — The Real Thing," Cambridge University Press, 2004.

[21] Kahneman, D., "Thinking, fast and slow," Allen Lane, 2011.

[22] Kooi, B. P., "Knowledge, Chance, and Change," Ph.D. thesis, Groningen University (2003).

[23] Kyburg, H., "Probability and the Logic of Rational Belief," Wesleyan University Press, Middletown, CT, 1961.

[24] Leitgeb, H., *The stability theory of belief*, Philosophical Review **123** (2014), pp. 131–171.

[25] Neumann, J. v. and O. Morgenstern, "Theory of Games and Economic Behavior," Princeton University Press, 1944.

[26] Plaza, J. A., *Logics of public communications*, in: M. L. Emrich, M. S. Pfeifer, M. Hadzikadic and Z. W. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, 1989, pp. 201–216.

[27] Ramsey, F., *Truth and probability*, in: R. Braithwaite, editor, *The Foundations of Mathematics and Other Essays*, Humanities Press, 1931 .

[28] Santoli, T., *Haskell project epistemic logic*, Technical report, ILLC (Summer 2014).

[29] Savage, L. J., "The Foundations of Statistics — Second Revised Edition," Dover, New York, 1972.

[30] Scott, D., *Measurement structures and linear inequalities*, Journal of Mathematical Psychology **1** (1964), pp. 233–247.

[31] Segerberg, K., *Qualitative probability in a modal setting*, in: J. Fenstad, editor, *Proceedings of the 2nd Scandinavian Logic Symposium* (1971), pp. 341–352.

[32] Talbott, W., *Bayesian epistemology*, in: E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, Stanford University, Fall 2013 Edition `http://plato.stanford.edu/archives/fall2013/entries/epistemology-bayesian/`.

[33] Walley, P. and T. Fine, *Varieties of modal (classificatory) and comparative probability*, Synthese **41** (1979), pp. 321–374.

[34] Zvesper, J. A., "Playing with Information," Ph.D. thesis, ILLC, University of Amsterdam (2010).