The logic of evidence and its relation to rational belief

October 31, 2014 *Seamus Bradley*

Evidence

Evidence is an important concept in formal epistemology, but it often takes on a subsidiary role to belief: evidence is (merely) what constrains rational belief. Evidence is then discussed in terms of the role it plays in informing change in rational belief. In contrast, I place evidence centre stage. I discuss the formal constraints on change in evidence base in the light of new information. Drawing on the AGM theory of belief revision, I discuss addition, contraction and revision operations for evidence bases. I extend this to the idea of probabilistic evidence bases.

Let $\mathbf{E}$ be a set of sentences of some formal language, and let $\tilde{\mathbf{E}}$ be the conjunction of those sentences. $\mathbf{E}$ is your EVIDENCE BASE and $\tilde{\mathbf{E}}$ is your EVIDENTIAL STATE. Following Williamson (forthcoming) we understand evidence as what is *granted* in a certain context.

Evidence informs belief

Using Isaac Levi's idea of a *Confirmational Commitment* we can use the idea of an evidence base to inform us about what our rational beliefs ought to be (Levi 1980).[1] For our purposes, a confirmational commitment (CC) is a function from evidence bases to probability functions. It takes, as input, your evidence at a time and outputs what your rational belief should be. SUBJECTIVIST CCs are quite permissive, whereas OBJECTIVIST CCs are much more constrained. An absolutely core idea of being responsive to evidence is that if $E \in \mathbf{E}$ is in your evidence, then you should fully believe E. The Subjective Bayesian CC is defined as follows:

$$C_{SB}(\mathbf{E}) = \mathbf{pr}(-|\tilde{\mathbf{E}})$$

As it stands, $\mathbf{pr}$ is an unconstrained prior probability function. We contrast this with the Objective Bayesian CC (Williamson 2010). First, let $\mathbf{E}^\dagger = \{\mathbf{pr} \in \mathbb{P}, \mathbf{pr}(\tilde{\mathbf{E}}) = 1\}$. That is, $\mathbf{E}^\dagger$ is the set of probability functions satisfying the constraint to fully believe the evidence.

$$C_{OB}(\mathbf{E}) = \Downarrow \mathbf{E}^\dagger$$

For set of probability functions $\mathbf{P}$, define $\Downarrow \mathbf{P}$ as the maximally equivocal member of $\mathbf{P}$ (Williamson 2010, section 3.4). No reference is made to a prior probability here.

Change in evidence

Let's say you obtain new information that A is true, and this information does not contradict any of your old evidence. If $\mathbf{E}$ was your old evidence, what should your new evidence be? Call this $\mathbf{E}^+_A$. We can use the literature on AGM belief revision (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988) as a starting point for discussion of the logic of evidence. Some interesting connections between graded belief change and full belief change have been drawn (Gärdenfors 1988; Makinson 2011). These connections are typically of the form "If your degrees of belief are updated in this way, then your associated full beliefs are that way". Our goal is slightly different. We want to understand the converse: if your evidence is this way, and evidence constrains graded belief in that way, how do your graded beliefs change? It's fairly easy to show that under fairly mild conditions, expansion – addition of evidence consistent with $\mathbf{E}$ – yields conditionalisation for the subjectivist and MaxEnt updating for the objectivist.

More interesting are the cases of contraction and revision. We discuss these cases, and relate contraction and revision to undermining evidence and the concept of a defeater. We also relate these ideas to admissibility in the context of some of chance-credence coordination norms (calibration, direct inference, principal principle etc).

We discuss the results of Lin and Kelly (2012)

[1] Levi actually introduces confirmational commitments in the context of criticising the sort of Bayesian views I discuss in what follows. Discussing this would take us too far afield.

who show that the AGM framework cannot accommodate revision that always tracks conditionalisation. This does not undermine using AGM as a theory of evidence, since you cannot consistently grant (endorse as part of your evidential state) all of the propositions involved in preface and lottery style problems. The scope of the conditionalisation norm (for subjectivist CCs) is clearly circumscribed.

Probabilistic evidence

Propositional evidence bases are somewhat limited. How can you accommodate natural kinds of evidential constraints like "X and Y are unrelated" or "Z is more likely than W"? Indeed, chance propositions seem like an awkward roundabout way of crowbarring chance constraints into a propositional framework. It's much more natural to consider an evidence base as a set of probabilistic constraints that your credence must satisfy. A natural extension of the subjectivist CC for probabilistic evidence is given by cross-entropy minimisation or similar (Diaconis and Zabell 1982). The objectivist CC obviously also naturally covers probabilistic evidence bases.

We shall also discuss revision of probabilistic evidence bases through viewing them as sets of sentences in a formal language that describe constraints on rational belief and consider an AGM-style theory of expansion, contraction and revision for sets of constraints.

Generic constraints and admissibility

We finish with a difficulty that is currently unresolved in this framework. Consider the constraint G: "this coin is fair". This seems to entail that for a particular (standard)[2] toss of the coin you should believe it will land heads up to degree 0.5; call this constraint P. We have a GENERIC chance claim G, and a PARTICULAR chance claim P. Imagine that we are considering an already completed toss of said coin, where you know that the coin landed tails. This is a clear example of INADMISSIBLE evidence, and on the view presented here, learning inadmissible evidence involves contracting the particular chance claim. We don't want such contraction to involve contracting the generic chance claim. But if the relation between the generic and particular constraints is one of entailment, then we *must* contract the generic claim. The only contraction postulates we need to endorse to generate this worry are that contraction is successful and that evidence is closed under logical consequence.

[2]The caveat "standard" is intended to rule out any funny business like a skilled magician who can toss a "fair" coin so as to almost always land heads.

References

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

Diaconis, P. and S. Zabell (1982). "Updating Subjective Probability". *Journal of the American Statistical Association* 77, pp. 822–830.

Gärdenfors, P. (1988). *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press.

Levi, I. (1980). *The Enterprise of Knowledge*. The MIT Press.

Lin, H. and K. T. Kelly (2012). "Propositional Reasoning that Tracks Probabilistic Reasoning". *Journal of Philosophical Logic* 41, pp. 957–981.

Makinson, D. (2011). "Conditional probability in the light of qualitative belief change". *Journal of Philosophical Logic* 40, pp. 121–153.

Williamson, J. (2010). *In Defense of Objective Bayesianism*. Oxford University Press.

——— (forthcoming). "Deliberation, Judgement and the Nature of Evidence". *Economics and Philosophy*.

Uncertain deduction with ifs, ands, and ors

Nicole Cruz[1,3], Jean Baratgin[3,4], Mike Oaksford[1], & David Over[2]
Birkbeck, University of London[1]; University of Durham[2]; EPHE[3]; Paris 8 University[4]

People's everyday reasoning, both in the context of science and on the streets, mostly concerns statements that they do not assume to be true or take to be certain, but which they only hold with varying degrees of belief. The psychology of deductive reasoning has recently started to study such inferences from uncertain beliefs. This new development is part of a paradigm shift from binary to Bayesian / probabilistic approaches in cognitive psychology(Evans & Over, 2013; Oaksford&Chater, 2013; Pfeifer &Kleiter, 2010). In the earlier binary approach to the study of deductive inference, participants in experiments were asked to assume that the premises were true.

A central developmentinthe probabilistic approach was to generalize binary validity and consistency to cover uncertain premises. Probabilistic notions of validity and consistency can be defined: *p-validity* and *coherence*. To define p-validity, let the uncertainty of a statement equal 1 minus its probability: $U(p) = 1 - P(p)$. Then an inference is p-valid if and only ifthe uncertainty of the conclusion cannot begreater than the sum of the uncertainties of the premises (Adams, 1998). A p-valid inference does not increase uncertainty. The concept of coherence is defined as respecting the axioms of probability theory. These are normative principles that people do not always conform to. For example, they sometimes commit the conjunction fallacy of judging the probability of a conjunction as higher than the probability of one of its conjuncts (Tversky&Kahneman, 1983).

Recent psychological research has provided strong support for the conditional probability hypothesis that people judge the probability of the conditional to be the conditional probability: $P(\text{if } p \text{ then } q) = P(q|p)$. A conditional that satisfies this relation has been called the probability conditional; its p-valid inferences differ from those of the material conditional. The question of whether people conform to p-validity and coherence for natural language conditionals has only just started to be investigated(Evans, Thompson, & Over, 2013, Pfeifer &Kleiter, 2005, 2010; Singmann, Klauer, & Over, 2014). People were found to respect p-validity and coherence at above chance levels for the p-valid inference MP, and to respect coherence at above chance levels for the inferences MP and DA. It was also found that the degree to which people take p-validity and coherence into account increases when they are given the explicit task of making these inferences(Evans, Thompson, & Over, 2013).

We conducted two experiments to study the extent to which people conform to p-validity and coherence for a novel set of inferences between conditionals, conjunctions, and disjunctions. We also assessed the effect of an explicit inference task. Six inferences were examined: (1) *p, therefore p or q*, (2) *If not-p then q, therefore p or q*, (3) *p or q, therefore if not-p then q*, (4) *p & q, therefore if p then q*, (5) *p, q, therefore if p then q*, and (6) *p& q, therefore p*.Experiment 1 concerned inferences (1) to (3) and was conducted online (n = 871). Experiment 2 concerned inferences (4) to (6) and was conducted in the lab (n = 48).Participants in both experiments were divided into two groups of approximately equal size. The inferencegroup evaluated explicit inferences, whereas the statements group evaluated thestatements from the inferences separated from each other. In both groups, participants were first presented with a short scenario describing a person. They were then shown a series of statements, or explicit inferences, about the person. For each statement, they were asked how much confidence they would have in the statement. For each explicit inference, they were asked how much confidence they would have in the premise(s) of the argument, and how muchin the conclusion, given the premise(s). Participants provided their answers in percent by writing a number between 0 and 100 in a box next to each statement. The experiments followed a mixed design with task (statements, inferences) as a between subjects variable and inference type as within subjects variable. Because Experiment 1 used statements containing negations, it varied in addition the position of the negation to control for negation effects.

Participants' responses werefound to be in accordance with p-validity and coherence at above chancelevels when the task was to draw explicit inferences, but not when the task was to evaluatethe statements from inferences separated from each other. An exception to this overall pattern was inference (6), which looked for the conjunction fallacy. Participants tended to commit the fallacy, giving responses at odds with p-validity and coherence in both the statements and the inference group.

Our results provide evidence that people respect normative criteria for uncertain deduction, at least when explicitly engaging in reasoning. An exception is the conjunction fallacy, which is confirmed to be a robust bias found even in explicit inferences. Further experiments could test if our findings can be extended to abstract materials, and could try to disentangle the contribution of task instructions and of working memory constraints on the difference found between the statements and the inference group.

References

Adams, E. (1998). A primer of probability logic.Stanford: CLSI Publications.

De Finetti, B. (1936/1995).The logic of probability.*Philosophical Studies, 77*(1)*,* 181-190.

Evans, J. S. B. T., & Over, D. E. (2013). Reasoning to and from belief: Deduction and induction are still distinct. *Thinking & Reasoning, 19*(3), 267-283.

Evans, J. St. B. T., Thompson, V., & Over, D. E. (2013).Uncertain deduction and the new psychology of conditional inference.University of Plymouth manuscript.

Oaksford, M., &Chater, N. (2013).Dynamic inference and everyday conditional reasoning in the new paradigm.*Thinking & Reasoning, 19*(3-4), 246-379.

Pfeifer, N. &Kleiter, G. D. (2005).Towards a mental probability logic.PsychologicaBelgica, 45 (1), 71_99.

Pfeifer, N. & Kleiter, G. D. (2010). The conditional in mental probability logic. In M.Oaksford & N. Chater (Eds.), *Cognition and conditionals: Probability and logic in human thinking* (pp. 153-173). Oxford: Oxford University Press.

Ramsey, F. P. (1926/1990). Truth and probability. In D. H. Mellor (Ed.), Philosophicalpapers (pp. 52–94). Cambridge: Cambridge University Press.

Singmann, H., Klauer, K. C., & Over, D. (2014).New normative standards of conditional reasoningand the dual-source model.*Frontiers in Psychology.*DOI: 10.3389/fpsyg.2014.00316.

Tversky, A., &Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90*(4), 293-315.

Reconditioning the Conditional

David Miller

http://www.warwick.ac.uk/go/dwmiller

Abstract

That there are connections between classical sentential logic and the elementary theory of probability is not open to serious dispute. In the formal theory of probability axiomatized in *The Logic of Scientific Discovery* (Popper 1959), appendix $*$v, it is possible to give a variety of rigorous definitions of the classical deducibility of one statement $\mathfrak{c}$, the *conclusion*, from another statement (or the same one) $\mathfrak{a}$, the *assumption* (or premise), a relation that will be written $\mathfrak{a} \vdash \mathfrak{c}$.

It would be a mistake to suppose, however, that probability is the only way, or even the best way, in which degrees of classical deducibility may be introduced. This paper will be concerned with an alternative measure, the measure of *deductive dependence* of one statement on another statement (or the same one), which was introduced by Miller & Popper (1986), and considered from other perspectives by Hempel & Oppenheim (1948), by Reichenbach (1954), and Hilpinen (1970). It will be shown that replacement of the probability function by the deductive dependence function sheds light on the *problem of indicative conditionals*, one of the most tenaciously unsolved problems of modern philosophical logic, and especially on the hypothesis of the *conditional construal of conditional probability* (facetiously dubbed CCCP by Hájek & Hall 1994).

For clarity we begin with an abstract probability function $\mathfrak{p}$ on the algebra of statements, and define the *credence* function $\mathfrak{c}(\mathfrak{c} \mid \mathfrak{a})$ by the identity $\mathfrak{c}(\mathfrak{c} \mid \mathfrak{a}) = \mathfrak{p}(\mathfrak{c} \mid \mathfrak{a})$. The *deductive dependence* function $\mathfrak{q}(\mathfrak{c} \mid \mathfrak{a})$, is defined as equal to $\mathfrak{p}(\mathfrak{a}' \mid \mathfrak{c}')$, where the prime is a symbol for classical negation. Since $\mathfrak{c}' \vdash \mathfrak{a}'$ if & only if $\mathfrak{a} \vdash \mathfrak{c}$, by contraposition, $\mathfrak{q}(\mathfrak{c} \mid \mathfrak{a})$, like $\mathfrak{c}(\mathfrak{c} \mid \mathfrak{a})$, takes the value 1 when $\mathfrak{a} \vdash \mathfrak{c}$. It is in this sense that the functions $\mathfrak{c}$ and $\mathfrak{q}$ each provides a generalization of the relation of deducibility. But the conditions under which $\mathfrak{c}$ and $\mathfrak{q}$ take the value 0 are very different. If $\mathfrak{a}$ and $\mathfrak{c}$ are mutual *contraries*, and $\mathfrak{a}$ is consistent, that is, if $\mathfrak{ac} \equiv \bot \not\equiv \mathfrak{a}$, then $\mathfrak{c}(\mathfrak{c} \mid \mathfrak{a}) = 0$. This condition is also necessary if the underlying abstract probability measure $\mathfrak{p}$ is regular. Equally, if $\mathfrak{a}$ and $\mathfrak{c}$ are mutual *subcontraries*, and $\mathfrak{c}$ is non-trivial, that is, if $\mathfrak{a} \vee \mathfrak{c} \equiv \top \not\equiv \mathfrak{c}$, then $\mathfrak{q}(\mathfrak{c} \mid \mathfrak{a}) = 0$. This condition too is necessary if the probability measure $\mathfrak{p}$ is regular. It follows that $\mathfrak{q}(\mathfrak{c} \mid \mathfrak{a})$ may exceed 0 when $\mathfrak{a}$ and $\mathfrak{c}$ are mutually inconsistent. Unless a positive degree of belief is possible in a conclusion $\mathfrak{c}$ in the presence of an assumption $\mathfrak{a}$ that contradicts it, the function $\mathfrak{q}$ is not a measure of degree of belief. What $\mathfrak{q}(\mathfrak{c} \mid \mathfrak{a})$ is is a measure of the extent to which (the content of) the assumption $\mathfrak{a}$ approximates (the content of) the conclusion $\mathfrak{c}$.

The CCCP hypothesis takes a variety of distinguishable forms, such as *Adams's thesis* and *Stalnaker's thesis*, which we need not here distinguish. We shall attend to one of its weaker formulations, which we call CCCP$_0$: $\forall\mathfrak{a}\forall\mathfrak{c}\exists\mathfrak{y}\forall\mathfrak{b}\, \mathfrak{c}_\mathfrak{b}(\mathfrak{y}) = \mathfrak{c}_\mathfrak{b}(\mathfrak{c} \mid \mathfrak{a})$. The intention is that for each $\mathfrak{a}, \mathfrak{c}$ the statement $\mathfrak{y}$ postulated by CCCP$_0$ can be identified with the indicative conditional *if* $\mathfrak{a}$ *then* $\mathfrak{c}$, and that its credence equals the credence of its consequent $\mathfrak{c}$ given its antecedent $\mathfrak{a}$, not only under the measure $\mathfrak{p}$ from which $\mathfrak{c}$ derives, but under any measure $\mathfrak{p}_\mathfrak{b}$ obtained from $\mathfrak{p}$ by Bayesian conditionalization on the statement $\mathfrak{b}$. Lewis (1976) and others have shown that CCCP$_0$ is vacuous: there is a statement $\mathfrak{y}$ satisfying it only if $\mathfrak{p}$ is two-valued. There have been copious subsequent attempts both to rescue CCCP and to extend and to intensify the arguments against it.

In the words of Arló-Costa (2001), nonetheless, it is a 'highly entrenched tenet of probabilistic semantics [that] the assertability of conditionals goes by conditional probability'. That is, whatever the agent's credence function $\mathfrak{c}$ may be, the assertability, or sometimes its acceptability (the

difference is not important here), of the indicative conditional $\mathfrak{a} \rightsquigarrow \mathfrak{c}$ is taken to be measured by the value of the credence $\mathfrak{c}(\mathfrak{a} \rightsquigarrow \mathfrak{c})$. When $\mathfrak{c}(\mathfrak{a} \rightsquigarrow \mathfrak{c}) = 1$, then the consequent $\mathfrak{c}$ is (almost) fully assertable on the basis of $\mathfrak{a}$, and when $\mathfrak{c}(\mathfrak{a} \rightsquigarrow \mathfrak{c}) = 0$, then $\mathfrak{c}$ is (almost) fully deniable on the basis of $\mathfrak{a}$. It is evident that the assertability of *if* $\mathfrak{a}$ *then* $\mathfrak{c}$ is understood to be a measure of *how justified the agent is* in asserting $\mathfrak{c}$ on the basis of the information or evidence $\mathfrak{a}$. Locutions such as *probably*, and *in my opinion*, and *I think*, are commonly used to qualify statements that are not fully asserted. The less probable that $\mathfrak{c}$ is, given $\mathfrak{a}$, the less the agent is entitled to assert it, or the more tentatively he asserts it. In this vein, Lucas (1970) called probability 'a guarded guide'.

Those of us who take seriously the goal of truth, but dismiss as not quite serious the goal of justified truth, do not worry whether we are ever entitled to assert a statement. We think that we are entitled to say what we like, whatever epistemological authoritarians may enjoin. But if we know that an asserted statement is not true, we may qualify our assertion by such expressions as *about* or *or so* or *roughly* or *more or less.* Since the quantity $\mathfrak{q}(\mathfrak{c} \,|\, \mathfrak{a})$, the deductive dependence of a non-tautological statement $\mathfrak{c}$ on a statement $\mathfrak{a}$, is a straightforward measure of how well (the content of) $\mathfrak{c}$ is approximated by (the content of) $\mathfrak{a}$, ranging from 0, when $\mathfrak{a}$ contains none of $\mathfrak{c}$, to 1 when it contains it all, it does appear that $\mathfrak{q}(\mathfrak{c} \,|\, \mathfrak{a})$ may serve as another kind of measure of the assertability or the acceptability of the statement $\mathfrak{c}$ in the presence of $\mathfrak{a}$. If our aim is truth, then the higher $\mathfrak{q}(\mathfrak{c} \,|\, \mathfrak{a})$ is, the more successful is the statement (or hypothesis $\mathfrak{c}$), given the statement (or evidence) $\mathfrak{a}$. It is a pleasant surprise that $\mathfrak{q}_{\mathfrak{b}}(\mathfrak{c} \,|\, \mathfrak{a}) = \mathfrak{q}(\mathfrak{b} \to \mathfrak{c} \,|\, \mathfrak{a})$, where $\to$ is the usual material conditional. In short, if the credence function $\mathfrak{c}$ is replaced by the deductive dependence function $\mathfrak{q}$, then an appropriately modified version of CCCP not only avoids vacuity but is demonstrably true. Lewis's results, disastrous for $\mathfrak{p}$, are smoothly avoided by $\mathfrak{q}$.

The crucial difference between the function $\mathfrak{c}$ and $\mathfrak{q}$ is that for $\mathfrak{c}$, but not for $\mathfrak{q}$, the process of Bayesian *conditionalization*, the generally agreed way in which a probability distribution is updated on the receipt of new information, has the same mathematical effect as the application of the probability functor $\mathfrak{p}$ not to a single argument (in the present paper, a statement) but to two arguments, or to one statement relative to another, yielding a binary measure $\mathfrak{p}(\mathfrak{c} \,|\, \mathfrak{a})$ that is standardly called *conditional probability.* These processes of *updating* and *relativization*, as they are better called, are interchangeable for $\mathfrak{p}$ and therefore for $\mathfrak{c}$. But for the function $\mathfrak{q}$, updating and relativization are distinct processes. This allows us understand how and why the replacement in the CCCP hypothesis of $\mathfrak{c}$ by $\mathfrak{q}$ makes such a dramatic difference.

The formula $\forall\mathfrak{a}\forall\mathfrak{c}\exists\mathfrak{y}\forall\mathfrak{b}\, \mathfrak{c}(\mathfrak{y} \,|\, \mathfrak{b}) = \mathfrak{c}(\mathfrak{c} \,|\, \mathfrak{a}\mathfrak{b})$ is a notational variant of CCCP_0 obtained by writing $\mathfrak{c}(\mathfrak{c} \,|\, \mathfrak{a}\mathfrak{b})$ for $\mathfrak{c}_{\mathfrak{b}}(\mathfrak{c} \,|\, \mathfrak{a})$. The formula CCCP_1 below is obtained from it by first commuting the terms in the conjunction $\mathfrak{a}\mathfrak{b}$, then interchanging the letters $\mathfrak{a}$ and $\mathfrak{b}$ throughout, and finally changing $\mathfrak{c}(\mathfrak{c} \,|\, \mathfrak{a}\mathfrak{b})$ back to $\mathfrak{c}_{\mathfrak{b}}(\mathfrak{c} \,|\, \mathfrak{a})$. It is because updating and relativization are interchangeable manoeuvres that CCCP_0 and CCCP_1 are logically equivalent, though they look rather different.

$$\text{CCCP}_0 \qquad \forall\mathfrak{a}\forall\mathfrak{c}\exists\mathfrak{y}\forall\mathfrak{b}\, \mathfrak{c}(\mathfrak{y} \,|\, \mathfrak{b}) = \mathfrak{c}_{\mathfrak{b}}(\mathfrak{c} \,|\, \mathfrak{a}),$$
$$\text{CCCP}_1 \qquad \forall\mathfrak{b}\forall\mathfrak{c}\exists\mathfrak{y}\forall\mathfrak{a}\, \mathfrak{c}(\mathfrak{y} \,|\, \mathfrak{a}) = \mathfrak{c}_{\mathfrak{b}}(\mathfrak{c} \,|\, \mathfrak{a}).$$

These formulas are both vacuous. But if we replace $\mathfrak{c}$ by $\mathfrak{q}$ in both CCCP_0 and CCCP_1, yielding

$$\text{CCCQ}_0 \qquad \forall\mathfrak{a}\forall\mathfrak{c}\exists\mathfrak{y}\forall\mathfrak{b}\, \mathfrak{q}(\mathfrak{y} \,|\, \mathfrak{b}) = \mathfrak{q}_{\mathfrak{b}}(\mathfrak{c} \,|\, \mathfrak{a}),$$
$$\text{CCCQ}_1 \qquad \forall\mathfrak{b}\forall\mathfrak{c}\exists\mathfrak{y}\forall\mathfrak{a}\, \mathfrak{q}(\mathfrak{y} \,|\, \mathfrak{a}) = \mathfrak{q}_{\mathfrak{b}}(\mathfrak{c} \,|\, \mathfrak{a}),$$

one of the resulting formulas is vacuous, but not the other. Putting $\top$ for $\mathfrak{b}$ in CCCQ_0 shows that the function $\mathfrak{q}$ must be two-valued. In contrast, CCCQ_1 holds non-trivially for all $\mathfrak{b}, \mathfrak{c}$, since $\mathfrak{y}$ may be the material conditional $\mathfrak{b} \to \mathfrak{c}$. As already noted, $\forall\mathfrak{b}\forall\mathfrak{c}\forall\mathfrak{a}\, \mathfrak{q}(\mathfrak{b} \to \mathfrak{c} \,|\, \mathfrak{a}) = \mathfrak{q}_{\mathfrak{b}}(\mathfrak{c} \,|\, \mathfrak{a})$. It will be argued that CCCQ_1, no less than CCCP_0, heeds the advice embodied in Ramsey's test.

Qualitative Probabilistic Inference with Default Inheritance for Exceptional Subclasses

Paul D. Thorn[*] and **Christian Eichhorn**[†] and **Gabriele Kern-Isberner**[‡] and **Gerhard Schurz**[§]
[†,‡]Technische Universität Dortmund [*,§]Heinrich Heine Universität Düsseldorf

There are qualitative systems of conditional reasoning (such as Adams' (1975) system P (cf. Kraus et al., 1990)) that can be used to make inferences about conditional probabilities, where conditionals of the form $\phi \Rightarrow \psi$ are treated as expressing that the conditional probability of ψ given ϕ is high. The interpretation of conditionals as expressing high conditional probabilities also provides an appropriate intuitive basis for evaluating reasoning systems where conditionals are understood as expressing inference rules whose application is appropriate by default (or defeasibly). For example, the conditional, $B \Rightarrow F$, may express that *birds are normally capable of flight*, and that, by default, it is appropriate to infer, for any particular bird, c, that c is capable of flight.

Within a system where conditionals are treated as expressing defaults, it is desirable that *subclass inheritance* among defaults be licensed, defeasibly. For example, from the default $B \Rightarrow F$, we would like to infer that $B \wedge C \Rightarrow F$ (crows normally fly), in the case where we have no background knowledge indicating that crows (a subclass of birds) are exceptional birds. Such inferences are assumed to be defeasible, meaning that there are conditions under which such inferences are defeated (i.e., conditions under which the inference is not licensed). It should be noted here that inferences of the present sort ($B \Rightarrow F$ / $B \wedge C \Rightarrow F$) are invalid in the case where conditionals are treated as expressing that a corresponding conditional probability is high. We should not regard this as a problem, in itself. Rather our concern to permit inferences via subclass inheritance, when combined with a probabilistic interpretation of conditionals, indicates our interest in developing systems of conditional reasoning that are inductive, in the broad sense.

Beyond defeasible subclass inheritance, it is controversial whether *inheritance in the case of exceptional subclasses* should be licensed, defeasibly (cf. Geffner and Pearl, 1992; Bacchus et al., 1996). For example, notice that penguins are exceptional birds inasmuch as they lack the capacity of flight. Given that penguins represent an exceptional subclass of the class of birds, it is controversial whether the subclass, penguins, should inherit other characteristics typical of birds. For example, assuming $B \Rightarrow W$ (birds generally have wings), is it reasonable to infer the default $B \wedge P \Rightarrow W$ (penguins normally have wings), given $B \wedge P \Rightarrow \neg F$ (penguins are normally not capable of flight)? In addition to presenting some informal considerations that favor of inheritance in the case of exceptional subclasses, we employ the methods of (Schurz and Thorn, 2012), in order to compare the reasonableness of systems of conditional reasoning that differ in whether or not they license default inheritance for exceptional subclasses. In particular, we compare the performance of system Z (Pearl, 1990), which does not license default inheritance for exceptional subclasses, with inference by c-representations (Kern-Isberner, 2001), which is hereafter referred to as system CR, which does (cf. Geffner and Pearl, 1992; Goldszmidt and Pearl, 1996; Kern-Isberner and Eichhorn, 2014).

Schurz and Thorn (2012) used computer simulations of a 'task environment' in order to evaluate four well known systems of conditional reasoning, known as O, P, Z, and QC (Hawthorne, 1996; Hawthorne and Makinson, 2007; Paris and Simmonds, 2009; Adams, 1975; Kraus et al., 1990; Suppes, 1966). The simulations proceeded by four basic steps. First, a probability distribution was generated at random. This probability distribution was thereafter understood as representing the true (stochastic) state of the world. Second, several true probability statements (given the true state of the world) were selected at random and passed to the four reasoning systems. Next, it was determined which conclusions were inferable by each system, given the statements provided. For systems P, Z, and QC, inferred probability bounds for derived conditionals were assigned using the improbability sum calculation which is valid for system P (see Schurz and Thorn, 2012, p. 584). In the case of system O, the inferred probability bounds for derived conditionals were set to be the probability of the least probable premise needed in inferring the conditional. Finally, each system was assigned numeric scores, according to the accuracy and informativeness of the conclusions that it drew. The preceding process was repeated thousands of times (ensuring small standard errors for the mean accuracy scores for the four systems), with variations of several key parameters, including: (1) the number of premise conditionals passed to the reasoning systems, (2) the minimum probability of these premises, and (3) the entropy level of the probability distribution that was understood as representing the true (stochastic) state of the world (Thorn and Schurz, 2013, 2014).

The four systems evaluated by Schurz and Thorn (2012) can be ordered according to the number of inferences that they license: O $\subset$ P $\subset$ Z $\subset$ QC. By appeal to the results of their simulations, Schurz and Thorn make a relatively convincing case that system Z provides a good balance of reward versus risk, in supporting accurate and informative conclusions and avoiding inaccurate ones. In other words, in considering (from left to right) the continuum of conditional reasoning systems represented by systems O, P, Z, and QC, system Z represents a good 'stopping point' regarding the set of conditional inferences one should accept: Limiting oneself to the inference licensed by system P would be too cautious, whereas performing all of the inferences licensed by system QC would be reckless.

While system Z licenses subclass inheritance (defeasibly), it does not license inheritance in the case of exceptional subclasses. Moving beyond the study of (Schurz and Thorn, 2012), we pursued a comparative study of the performance of system Z and system CR. Our preliminary results include the following:

1. While neither system Z nor system CR strictly includes the other (Kern-Isberner and Eichhorn, 2014), the set of system Z inferences was a subset of the set of system CR inferences within a vast majority of our simulations.
2. System CR licensed significantly more inferences than system Z, with a greater difference in the number of inferences made in cases where the entropy of the underlying probability distribution (representing the environment) was small.
3. The performance of both systems Z and CR is best when the entropy of the underlying probability distribution is high, with relatively poor performance when the entropy level is low (with a value less than 1.5 on a scale from 0 to 4).
4. The accuracy of inferences licensed by system CR was (approximately) identical to the accuracy of inferences drawn by system Z, so long as the entropy of the underlying probability distribution was moderate or high (at least 2.0 on a scale from 0 to 4).
5. Cases where system CR licensed inheritance for an exceptional subclass produced conclusions whose accuracy was (approximately) identical to the accuracy of those conclusions drawn by system Z that were not drawn by system P.

Our results suggest that inference by system CR (and inheritance for exceptional subclasses, as licensed by system CR) is *as reasonable* as inference by system Z. Our results also show that for practical purposes, system CR represents a stronger system of inference than system Z. So in considering the continuum of conditional reasoning systems represented by systems O, P, Z, CR, and QC, system CR represent a new good stopping point in considering the set of conditional inferences one should perform.

References

Adams, E. (1975). *The Logic of Conditionals: An Application of Probability to Deductive Logic*. Synthese Library. Springer.

Bacchus, F., Grove, A. J., Halpern, J. Y., and Koller, D. (1996). From statistical knowledge bases to degrees of belief. *Artificial Intelligence*, 87(1–2):75–143.

Geffner, H. and Pearl, J. (1992). Conditional entailment: Bridging two approaches to default reasoning. *Artificial Intelligence*, 53(2–3):209–244.

Goldszmidt, M. and Pearl, J. (1996). Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 84(1–2):57–112.

Hawthorne, J. (1996). On the logic of non-monotonic conditionals and conditional probabilities. *Journal of Philosophical Logic*, 25:185–218.

Hawthorne, J. and Makinson, D. (2007). The quantitative/qualitative watershed for rules of uncertain inference. *Studia Logica*, 86(2):247–297.

Kern-Isberner, G. (2001). *Conditionals in Nonmonotonic Reasoning and Belief Revision – Considering Conditionals as Agents*. Number 2087 in LNCS. Springer.

Kern-Isberner, G. and Eichhorn, C. (2014). Structural inference from conditional knowledge bases. *Studia Logica*, 102(4):751–769.

Kraus, S., Lehmann, D., and Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207.

Paris, J. B. and Simmonds, R. (2009). O is not enough. *The Review of Symbolic Logic*, 2:298–309.

Pearl, J. (1990). System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *Proceedings of the 3rd conference on Theoretical aspects of reasoning about knowledge*, TARK '90, pages 121–135, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

Schurz, G. and Thorn, P. D. (2012). Reward versus risk in uncertain inference: Theorems and simulations. *The Review of Symbolic Logic*, 5:574–612.

Suppes, P. (1966). Probabilistic inference and the concept of total evidence. In Hintikka, J. and Suppes, P., editors, *Aspects of Inductive Logic*, pages 49–65.

Thorn, P. D. and Schurz, G. (2013). Ampliative Inference Under Varied Entropy Levels. In Beierle, C. and Kern-Isberner, G., editors, *Proceedings of th 4th Workshop on Dynamics of Knowledge and Belief (DBK-2013)*, pages 46–60. FernUniversität in Hagen.

Thorn, P. D. and Schurz, G. (2014). A Utility Based Evaluation of Logico-probabilistic Systems. *Studia Logica*, 102(4):867–890.

Borderlines, Borderline Probabilities and Probabilities of Borderlines

Jonathan Lawry, Department of Engineering Mathematics, University of Bristol, BS8 1UB, UK

There is a highly interconnected relationship between vagueness and uncertainty. It is not just that vagueness occurs in conjunction with epistemic uncertainty but also that linguistic uncertainty is integral to vague propositions themselves. The latter refers to uncertainty about the definition or interpretation of concepts in natural language and is a natural result of the empirical manner in which language is learnt. Lawry [6] and Lassiter [5] argue that this form of uncertainty is epistemic in nature and can be modelled probabilistically. In this case, the blurred boundary of a vague category can be modelled by probability defined over possible precise boundaries. There is nonetheless an important distinction between blurred boundaries and the explicit identification of borderline cases. Indeed the latter does not refer to epistemic uncertainty at all but instead results from a non-Boolean truth model. For example, given an exact value for Ethel's height it might be certain that she is borderline short. Vagueness is not only the result of linguistic uncertainty or of borderline cases but comprises of at least both of these features. Furthermore, vague predicates are everywhere embedded in our statements and beliefs about the world. Consequently, to assess such beliefs we must consider vagueness in conjunction with epistemic uncertainty about the state of the world. This requires an integrated approach capturing both uncertainty about the world and linguistic uncertainty about the conventions of language, together with non-Boolean truth models resulting from more flexible category representation.

In this paper we investigate these ideas in a propositional logic setting by defining probabilities over three valued valuations, i.e. taking truth values true, borderline or false, so as to generate upper and lower belief measures on the sentences of the language. Initially, we adopt an axiomatic approach and consider what properties should be satisfied by three valued valuations if they are to appropriately represent explicitly borderline cases, and following on from this we then investigate the relationship between two different types of valuations. We argue that the notion of *borderline* cannot be satisfactorily defined in terms of intermediate probability values. Finally, we propose belief pairs corresponding to lower and upper measures generated from probability distributions defined over three valued valuations, and discuss their properties.

Let $\mathcal{L}$ be a language of propositional logic with connectives $\wedge$, $\vee$ and $\neg$ and propositional variables $\mathcal{P} = \{p_1, \ldots, p_n\}$. Let $S\mathcal{L}$ denote the sentences of $\mathcal{L}$ as generated recursively from the propositional variables by application of the three connectives. The general definition of a three valued valuation on $\mathcal{L}$ is then given as follows:

Definition 1. Three Valued Valuation: *A three valued valuation on $\mathcal{L}$ is a function $\mathbf{v} : S\mathcal{L} \to \{1, \frac{1}{2}, 0\}$ such that $\forall \theta, \varphi \in S\mathcal{L}$ if $\mathbf{v}(\theta) \in \{0,1\}$ and $\mathbf{v}(\varphi) \in \{0,1\}$ then $\mathbf{v}(\neg\theta) = 1 - \mathbf{v}(\theta)$, $\mathbf{v}(\theta \wedge \varphi) = \min(\mathbf{v}(\theta), \mathbf{v}(\varphi))$ and $\mathbf{v}(\theta \vee \varphi) = \max(\mathbf{v}(\theta), \mathbf{v}(\varphi))$. Here the truth values denote absolutely true (1), borderline ($\frac{1}{2}$) and absolutely false (0) respectively. The restriction on $\mathbf{v}$ is that it should obey the same rules as Tarski valuations in the case of Boolean expressions.*

Two well known examples of valuations of this form are supervaluations and Kleene valuations:

- *Supervaluations [2]:* Let $\mathbb{T}$ denote the set of Tarski (classical) valuations defined on $\mathcal{L}$. A supervaluation is a three valued valuation define by a set $\Pi \subseteq \mathbb{T}$ of Tarski valuations corresponding to admissible precisifications, such that $\forall \theta \in S\mathcal{L}$;

$$\mathbf{v}(\theta) = \begin{cases} 1 : \min\{v(\theta) : v \in \Pi\} = 1 \\ 0 : \max\{v(\theta) : v \in \Pi\} = 0 \\ \frac{1}{2} : \text{otherwise} \end{cases}$$

- *Kleene valuations [4]:* A Kleene valuation is a three valued valuation defined recursively such that $\forall \theta, \varphi \in S\mathcal{L}$; $\mathbf{v}(\neg\theta) = 1 - \mathbf{v}(\theta), \mathbf{v}(\theta \wedge \varphi) = \min(\mathbf{v}(\theta), \mathbf{v}(\varphi))$ and $\mathbf{v}(\theta \vee \varphi) = \max(\mathbf{v}(\theta), \mathbf{v}(\varphi))$

In this context, it is also tempting to try to associate the borderline truth value of $\frac{1}{2}$ with intermediate probability values of an underlying probability measure on $S\mathcal{L}$. More formally, let $\rho : S\mathcal{L} \to [0,1]$ be a probability measure on $S\mathcal{L}$ and $\beta \in (0.5, 1]$ be a threshold value. We can then consider the mapping $\mathbf{v} : S\mathcal{L} \to \{0, \frac{1}{2}, 1\}$ such that $\forall \theta \in S\mathcal{L}$, $\mathbf{v}(\theta) = 1$ if $\rho(\theta) \geq \beta$, $\mathbf{v}(\theta) = 0$ if $\rho(\neg\theta) \geq \beta$ and $\mathbf{v}(\theta) = \frac{1}{2}$ if $1 - \beta < \rho(\theta) < \beta$. However, with the exception of the case in which $\beta = 1$ when we obtain a supervaluation, this approach does not generate a three valued valuation consistent with definition 1. This is because it is straightforward to pick ρ and β such that for some sentences θ and φ, $\mathbf{v}(\theta) = \mathbf{v}(\varphi) = 1$ whilst $\mathbf{v}(\theta \wedge \varphi) \neq 1$. This highlights the non-epistemic nature of the borderline truth value as given in definition 1 i.e. that it does not represent *uncertain* or *unknown*.

We will adopt an axiomatic approach so as to characterise both supervaluations and Kleene valuations in terms of intuitive properties on three valued valuations. In particular, supervaluations will be shown to be characterised by the following three axioms: $\forall \theta, \varphi \in S\mathcal{L}$;

- **P1 Duality:** $\mathbf{v}(\neg\theta) = 1 - \mathbf{v}(\theta)$.
- **P2 Tautology:** If $\models \theta$ then $\mathbf{v}(\theta) = 1$.
- **P3 Equivalence:** If $\theta \equiv \varphi$ then $\mathbf{v}(\theta) = \mathbf{v}(\varphi)$.

where $\models$ and $\equiv$ refer to the classical (Tarski) entailment and equivalence relations respectively. Alternatively, we will show that Kleene valuations are characterised by P1 together with the following additional axioms:

- **P4 Commutativity:** $\mathbf{v}(\theta \wedge \varphi) = \mathbf{v}(\varphi \wedge \theta)$ and $\mathbf{v}(\theta \vee \varphi) = \mathbf{v}(\varphi \vee \theta)$.
- **P5 Bounds:** If $\mathbf{v}(\theta) \neq 1$ or $\mathbf{v}(\varphi) \neq 1$ then $\mathbf{v}(\theta \wedge \varphi) \neq 1$, and if $\mathbf{v}(\theta) \neq 0$ or $\mathbf{v}(\varphi) \neq 0$ then $\mathbf{v}(\theta \vee \varphi) \neq 0$.
- **P6 Monotonicity:** If $\mathbf{v}(\psi) < \mathbf{v}(\varphi)$ then $\mathbf{v}(\theta \wedge \psi) \leq \mathbf{v}(\theta \wedge \varphi)$ and $\mathbf{v}(\theta \vee \psi) \leq \mathbf{v}(\theta \vee \varphi)$.
- **P7 Borderline:** If $\mathbf{v}(\theta) = \mathbf{v}(\varphi) = \frac{1}{2}$ then $\mathbf{v}(\theta \wedge \varphi) = \mathbf{v}(\theta \vee \varphi) = \frac{1}{2}$.

Focussing on the relationship between these two types of valuations we will then show that Kleene valuations are closely related to *complete bounded supervaluations* defined as follows: Let $\trianglelefteq$ be a partial ordering on $\mathbb{T}$ according to which $v_1 \trianglelefteq v_2$ if and only if $\forall p_i \in \mathcal{P}$, $v_1(p_1) \leq v_2(p_i)$. Then a complete bounded supervaluation is a supervaluation with the set of admissible precisifications of the form $\Pi = \{v \in \mathbb{T} : v_* \trianglelefteq v \trianglelefteq v^*\}$ where $\forall p_i \in \mathcal{P}$, $v_*(p_i) = \min\{v(p_i) : v \in \Pi\}$ and $v^*(p_i) = \max\{v(p_i) : v \in \Pi\}$.

Within the proposed framework, uncertainty concerning the sentences of $\mathcal{L}$ effectively corresponds to uncertainty as to which is the correct three valued valuation on $\mathcal{L}$. As outlined earlier we view uncertainty as being epistemic in nature, resulting from a lack of knowledge concerning either, the state of the world to which propositions refer, or the underlying definitions of concepts used in propositions. In the following we assume that this uncertainty is quantified by a probability distribution w defined over a finite set of three valued valuations $\mathbb{V}$.

Definition 2. Belief Pairs*: Let $\mathbb{V}$ be a finite set of three valued valuations and w be a probability distribution on $\mathbb{V}$ then we define a belief pair as a pair of lower and upper measures $\vec{\mu} = (\underline{\mu}, \overline{\mu})$ where $\underline{\mu}, \overline{\mu} : S\mathcal{L} \to [0,1]$ such that $\forall \theta \in S\mathcal{L}$; $\underline{\mu}(\theta) = w(\{\mathbf{v} \in \mathbb{V} : \mathbf{v}(\theta) = 1\})$ and $\overline{\mu}(\theta) = w(\{\mathbf{v} \in \mathbb{V} : \mathbf{v}(\theta) \neq 0\})$.*

In the cases that $\mathbb{V}$ is restricted only to supervaluations or to Kleene valuations we refer to $\vec{\mu}$ as a *supervaluation belief pair* or a *Kleene belief pair* respectively. Furthermore, $\vec{\mu}$ is a *complete bounded supervaluation belief pair* if $\mathbb{V}$ is restricted to complete bounded supervaluations. It is well known that *supervaluation belief pairs* correspond to Dempster-Shafer belief and plausibility measures on $S\mathcal{L}$ [3], [1]. However, the specific properties of complete bounded supervaluation belief pairs have only recently been studied and we recall some of them [8]. Kleene belief pairs have been proposed independently in [7] and [10]. We recap on their properties including a surprising characterisation of min-max fuzzy logic as shown in [9]. Finally, we extend a result in [8] so as to clarify the relationship between Kleene belief pairs and complete bounded supervaluation belief pairs as follows:

Definition 3. A Restricted Set of Sentences*: Let $\mathbb{A} = \{A \subseteq L\mathcal{L} : \forall p_i \in \mathcal{P}, \{p_i, \neg p_i\} \not\subseteq A, \{p_i, \neg p_i\} \cap A \neq \emptyset\}$ where $L\mathcal{L}$ denotes the literals of $\mathcal{L}$. For $A \in \mathbb{A}$, let $S\mathcal{L}_A \subseteq S\mathcal{L}$ denote the set of sentences of $\mathcal{L}$ generated recursively from A using only the connectives $\wedge$ and $\vee$. Then we define $S\mathcal{L}^* = \bigcup_{A \in \mathbb{A}} S\mathcal{L}_A$. Notice that $S\mathcal{L}^*$ is the subset of the sentences of $\mathcal{L}$ in negated normal form, for which it is not the case that both a propositional variable and its negation appear.*

Theorem 4. *Let $\vec{\mu}_1$ be a complete bounded supervaluation belief pair on $S\mathcal{L}$, then there is a Kleene belief pair $\vec{\mu}_2$ on $S\mathcal{L}$ such that $\forall \theta \in S\mathcal{L}^*$, $\vec{\mu}_1(\theta) = \vec{\mu}_2(\theta)$ and $\forall \theta \in S\mathcal{L}$, $\underline{\mu}_1(\theta) \geq \underline{\mu}_2(\theta)$ and $\overline{\mu}_1(\theta) \leq \overline{\mu}_2(\theta)$.*

References

[1] H. Field, (2000), 'Indeterminacy, Degree of Belief, and Excluded Middle', *Noûs*, Vol. 34, No. 1, pp1-30.

[2] K. Fine, (1975), 'Vagueness, Truth and Logic', *Synthese*, Vol. 30, pp265-300.

[3] J-Y. Jaffray, (1989), 'Coherent Bets Under Partially Resolving Uncertainty and Belief Functions', *Theory and Decision*, Vol. 26, pp99-105.

[4] S.C. Kleene, (1952), *Introduction to Metamathematics*, D. Van Nostrand Company Inc., Princeton, New Jersey.

[5] D. Lassiter, (2011), 'Vagueness as Probabilistic Linguistic Knowledge', *Lecture Notes in Artificial Intelligence*, Vol. 6517, pp127-150.

[6] J. Lawry, (2008), 'Appropriateness Measures: An Uncertainty Model for Vague Concepts', *Synthese*, Vol. 161(2), pp255-269.

[7] J. Lawry, I. Gonzalez-Rodriguez, (2011), 'A Bipolar Model of Assertabity and Belief', *International Journal of Approximate Reasoning*, Vol. 52, pp76-91.

[8] J. Lawry, Y. Tang, (2012), 'On Truth-gaps, Bipolar Belief and the Assertability of Vague Propositions', *Artificial Intelligence*, Vol. 191-192, pp20-41.

[9] J. Lawry, (2014), 'Probability, Fuzziness and Borderline Cases, *International Journal of Approximate Reasoning*, Vol. 55, pp1164-1184.

[10] J. Robert G. Williams, (2013), 'Probability and Non-classical Logic', forthcoming in (Eds. Hajek and Hitchcock), *Oxford Companion to the Philosophy of Probability*, OUP.

DEGREES OF TRUTH AS OBJECTIVE PROBABILITIES

Many-valued generalisations of classical (propositional) logic arise very naturally by extending the range of valuations beyond the binary set. Indeed since the mid 1950s, the extensive use of algebraic and functional-analytic techniques led to the development of the now flourishing field of mathematical fuzzy logic (see, e.g. Cintula et al., 2011), of which Łukasiewicz real-valued logic constitutes the leading edge (see e.g. Mundici, 2011). However, our firm grasp of the relevant mathematics is not matched by our understanding of the graded notion of truth that arises within such logics. In spite of the many attempts to link it to the formalisation of vagueness, the very notion of "graded truth" remains somewhat clouded in conceptual mystery.

This paper seeks to take some of that mystery out by putting forward a framework in which, under suitable conditions, degrees of truth are interpreted as objective probabilities. A central step in doing so consists in contrasting the *formal analogies* which hold between Łukasiewicz valuations and probability values with their *conceptual differences*. In order to spell out the latter we will rely on an axiomatisation of the relation $\succsim_T \subseteq \mathcal{SL}^2$ which we interpret as "θ is no less true than ϕ" (where as usual $\theta, \phi \in \mathcal{SL}$ are sentences built up by means of the classical connectives from a propositional language $\mathcal{L}$). As we argue in a companion paper, a suitable axiomatisation of the relation $\succsim_T$ proves to be sufficient for the existence of a Łukasiewicz valuation on $\mathcal{SL}$, and hence can be seen to provide an ordinal foundation to the latter. More precisely a small set of desirable properties on $\succsim_T$ guarantee the existence of a $v\colon \mathcal{SL} \to [0, 1]$ such that

$$\theta \succsim_T \phi \Longrightarrow v(\theta) \geq v(\phi)$$

where v is a Łukasiewicz valuation. Now, it is on this ordinal level of analysis that the contrast between degrees of truth and degrees of belief is best articulated, for this setting allows us to compare the axiomatisation of $\succsim_T$ with the classical axiomatisation of the relation "no less probable then", which can be shown to represent subjective probability functions (de Finetti, 1931; Savage, 1972).

Whilst the above adds a novel argument in support of the traditional distinction between degrees of truth and degrees of belief, the framework which underpins it also suggests the conditions under which the distinction vanishes. Hence our framework also accounts for the substantial formal overlap observed between the two notions. This, we argue, can be justified by interpreting degrees of truth as objective probabilities.

However the notion of objective probability is itself somewhat elusive. The concluding part of our paper is devoted to investigating two distinct instantiations of our framework which reflect two accounts of the very notion of objective probability. The first takes objective probability to arise as the "intersubjective agreement" of subjective probabilities. We will argue that this view supports an interpretation of degrees of truth as *ultimate degrees of belief* and compare this to the framework of Objective Bayesian Epistemology (Williamson, 2010). In our the second instantiation we will interpret objective probabilities as *chances* or physical probabilities (see e.g. Beisbart and Hartmann, 2011). Under this interpretation, we argue that the degree of truth of a sentence is best viewed as the chance that the event described by the sentence will obtain. Under both interpretations of objective probabilities, we will conclude, taking degrees of truth as objective probabilities improves our conceptual understanding of both notions.

References

C. Beisbart and S. Hartmann, editors. *Probabilities in Physics.* Oxford University Press, 2011.

P. Cintula, P. Hajek, and C. Noguera, editors. *Handbook of Mathematical Fuzzy Logic. Volumes 1-3.* College Publications, London, 2011.

B. de Finetti. Sul significato soggettivo della probabilità. *Fundamenta Mathematicae*, 17:289–329, 1931.

D. Mundici. *Advanced Lukasiewicz calculus and MV-algebras.* Springer, 2011.

L.J. Savage. *The Foundations of Statistics.* Dover, 2nd edition, 1972.

J. Williamson. *In defence of objective Bayesianism.* Oxford University Press, 2010.

ALGORITHMIC PROBABILITY AS OBJECTIVE-LOGICAL PRIOR DISTRIBUTION

Inspired by R. Carnap's *Logical Foundations of Probability*, R.J. Solomonoff sought to establish an objective-logical measure of the degree of confirmation $c(b, S)$ that a sequence S of symbols that we have witnessed confers upon the hypothesis that the symbol b will appear next. Ideas from information theory and computability theory led him to propose in the early sixties the definition of *algorithmic probability*. This definition has since often been presented as a "universal *a priori* distribution".

Solomonoff's proposal marks the birth of the theory of Kolmogorov complexity, and it can be seen as the progenitor to multiple successful modern approaches in statistics and machine learning (among them *prediction with expert advice* and the principle of *Minimum Description Length*). However, while numerous substantive claims about the theory's philosophical merits have been advanced from the angle of theoretical computer science, attention in the philosophical literature has so far been largely restricted to the occasional and uncritical mention in overview works. The aim of this talk is to sketch a more careful philosophical appraisal of Solomonoff's proposal, where we take our cue from Solomonoff's obligation to Carnap's programme of inductive logic. The driving question is whether, and in what sense, algorithmic probability can indeed be called a "universal" or "objective-logical" prior distribution.

It has often been noted that algorithmic probability retains an element of arbitrariness or subjectivity in the definition's dependence on the choice of a particular universal computer (Turing machine). After reviewing some arguments for the definition's objectivity, we will look at a key result that sheds new light on this discussion. This result is a Representation Theorem, akin to the famous theorem of de Finetti. It establishes that algorithmic probability is precisely a Bayesian mixture over the class of effective (computably approximable) hypotheses, whence the choice of universal Turing machine is precisely the choice of prior over this hypothesis class.

We can treat the algorithmic probability distribution as a prediction rule that conditionalizes on the incoming data. The Representation Theorem shows that we may then interpret algorithmic probability as a prediction rule adhering to a particular inductive assumption, or assumption on the patterns in the data that we may possibly encounter. This inductive assumption is the assumption of effectiveness (computable approximability) of data-generating sources.

The previous observation has a number of interesting consequences. As an important example, which will be briefly sketched: we can use it to challenge the prevalent feeling that algorithmic probability or Kolmogorov complexity gives an

Date: November 27, 2014.

objective measure of the simplicity of data sequences (and the further idea that we might even derive a justification for Occam's razor from it).

For the purpose of this talk, the main consequence of the above observation is that the question of the objectivity of algorithmic probability becomes the question of the unrestrictiveness (or universality) of the inductive assumption of effectiveness. Interestingly, while there is little ground for thinking that this inductive assumption is universal as an assumption on data-generating sources in the world, it can be argued more convincingly that it is universal as an assumption on competing prediction rules. If we accept this assumption that any competing prediction rule must be effective, then we can derive a number of formal results that establish that algorithmic probability as a prediction rule will predict well, if *any* predictor does.

In that sense, algorithmic probability possesses a universal optimality: the best we can do. Rather than in the tradition of Carnap, the original inspiration for Solomonoff, this view of Solomonoff's proposal is closer to a pragmatic approach to induction, in the spirit of Reichenbach.

Accuracy-Difficulty of a Single-Number Credence Representation in Belnap's Four-Valued Logic

An alternative philosophical approach to the representation of uncertain doxastic states is suggested, by considering how to model an agent who is concerned about the accuracy of her credences in Belnap's four-valued logic B_4. The paper has three parts. The first part motivates the whole project by showing that the single-number representation of credences is not appropriate for measuring an agent's accuracy in B_4. In the second part, we will introduce ordered pairs to represent an agent's uncertain doxastic states, and we will show how they can solve the problem faced by the single-number representation. In the third part, we will show how Joyce's idea of non-pragmatic vindication of probabilism works for ordered pairs in the classical and non-classical logical systems.

1. **Motivation**

Suppose that we want to measure, using the Brier score, the accuracy of an agent's credences in a proposition X in B_4. We will face one of the following consequences. The agent will either be a priori assigned different scores for what seem to be symmetric mistakes or we will risk a situation in which an agent has inaccurate credences, but the Brier score says that her credences are fully accurate.

For the sake of argument, suppose that an agent happens to live in the world with B_4. Suppose that we assign the truth-value 1 to the status true, the value 0 to the status false and the value 0.5 to the status neither true nor false. What about the fourth status: true and false? We can assign it either one of the three values, $\{1, 0, 0.5\}$, that we have already used or some new value from the interval $(0, 0.5) \cup (0.5, 1)$.

1. $(0, 0.5) \cup (0.5, 1)$

   Believing that X is true and false when it is true and believing that X is true and false when it is false seem, from the purely accuracy point of view, to be symmetric mistakes. They should be evaluated with the same score. It will, however, not be the case with the values from $(0, 0.5) \cup (0.5, 1)$. Suppose that we use the value 0.6. When an agent believes that X is true and false, but it is, in fact, false, then her Brier score will be 0.36. If she believes that X is true and false, but it is, in fact, true, then her Brier score will be 0.16. The Brier score should, however, reflect the symmetricity of an agent's mistake and not be a priori higher in one case and lower in the other. Moreover, there is no reason to choose one value from $(0, 0.5) \cup (0.5, 1)$ rather than another.

2. $\{1, 0, 0.5\}$

   Suppose that we use 0.5. Suppose that an agent believes that a proposition X is true and false, but X is neither true nor false. The agent is obviously inaccurate, yet the Brier score returns the value 0 which means that it reckons the agent to be fully accurate.

2. **Ordered Pairs**

We will use ordered pair (a_X, b_X) to represent an agent's uncertain doxastic state, where a is the degree to which an agent believes that a proposition X is true and b is the degree to which an agent believes that X is false. Such represenation will help us to motivate, epistemically, an assignment of truth-values in B_4 such that the issue will be resolved. Briefly, we will use pairs of numbers to represent the truth-values of propositions. For example, we will use the pair $(1, 0)$ to represent that a proposition is true and the pair $(0, 1)$ to represent that the proposition is false. We will then define connectives for ordered pairs standing for truth-values that will enable us to reconstruct B_4: 1.) Δ is the operation on ordered pairs (a_X, b_X) and (a_Y, b_Y) such that $(a_X, b_X)\Delta(a_Y, b_Y) := (min\{a_X, a_Y\}, max\{b_X, b_Y\})$. 2.) $\bigtriangledown$ is the operation on ordered pairs (a_X, b_X) and (a_Y, b_Y) such that $(a_X, b_X) \bigtriangledown (a_Y, b_Y) := (max\{a_X, a_Y\}, min\{b_X, b_Y\})$. 3.) $\bullet$ is the operation on an ordered pair (a_X, b_X) such that $\bullet(a_X, b_X) := (b_X, a_X)$. We can easily modify the Brier score to use ordered pairs such that we will be able to measure accuracy of an agent's credences in B_4 without further issues.

3. **Non-Pragmatic Vindication**

In the last part of the paper, we will show that Joyce's idea of non-pragmatic vindication of probabilism can be reproduced using ordered pairs not only for the classical two-valued logic, but also for non-classical systems. We will be mostly interested in the extensions into non-classical systems, especially, B_4. The challenging part is to provide a reasonable interpretation of elements of convex combinations for non-classical logical systems. The evaluation function from the second part that assigns ordered pairs as truth-values to propositions will serve as the interpretation of the weighted terms in convex combinations. To interpret the weights, we will need to identify a proposition not with a set of possible worlds where it is true, but with a pair of sets of possible worlds $(\{W_t\}, \{W_f\})$. The first set will be the set of possible worlds in which a proposition is true, and the second set will be the set of possible worlds in which a proposition is false. We will introduce a function P^* that will be a measure on the algebra of pairs of sets of possible worlds and we will show that $P^*(X)$ can be interpreted as an agent's credence in a proposition X i.e. in a pair of sets of possible worlds for X. We will axiomatize P^* and show that $B(\{W_t\}, \{W_f\})$ is in the convex hull (is not accuracy dominated) iff it follows the axiomatization of P^*, where B is some function from propositions i.e. $(\{W_t\}, \{W_f\})$ to ordered pairs of credences i.e. (a_X, b_X) for some X.

Scientific Reasoning, Bayesian Confirmation and Conditional Probabilities

November 5, 2014

Bayesianism is a theory of scientific reasoning phrased in terms of rational degrees of belief which are formalized through the probability calculus. Two cornerstones of Bayesianism are the principle that evidence E confirms hypothesis H if and only if $p(H|E) > p(H)$, and **Bayes' Theorem**:

$$p(H|E) = p(H)\frac{p(E|H)}{p(E)} \tag{1}$$

This mathematical result allows us to calculate the posterior probability $p(H|E)$ from the prior probability $p(H)$ and the likelihoods $p(E|H)$, $p(E|\neg H)$.

This approach works fine for assessing the confirmatory power of outcomes in future experiments. However, things are more complicated for assessing evidence from *past experiments*. Here is a simple example, the famous **Problem of Old Evidence**: when E is already known to the scientist, her prior degree of belief in E is maximal, and $p(E) = 1$. But with that assumption, it follows that the posterior probability of H cannot be greater than the prior probability: $p(H|E) = p(H) \cdot p(E|H) \leq p(H)$. Hence, E does not confirm H, contrary to our intuitions that past observations can strongly support a scientific theory.

The problem also arises for accounts of scientific confirmation that build on a comparison of the likelihoods of competing hypotheses (e.g., by means of Bayes factors) as long as probabilities are interpreted as rational degrees of belief. If my degree of belief in E is maximal because I know it has occurred in the past, also my degree of belief in E conditional on H, $\neg H$, etc. should be maximal. This seems to trivialize any confirmation theory based on a comparison of such likelihoods.

There are two main solution strategies:

1. A dynamic approach (Garber, Jeffrey, Niiniluoto, Earman, Sprenger), where we replace the question of whether old evidence E confirms H by the question of whether learning $H \vdash E$ confirms E. The latter question is relevant for reconstructing crucial episodes in the history of science, but it seems to dodge the principled challenge that the Problem of Old Evidence poses for Bayesian confirmation theory.

2. A counterfactual approach (Howson) where we assess the confirmation relation relative to *counterfactual degrees of belief*, namely the probability function where E has been removed from the agent's background assumptions K.

The counterfactual approach has been criticized for a number of reasons. One of them is that the probability function $p_{K-\{E\}}$ may not be well defined. However, while I am skeptic about

Howson's specific proposal, I believe it is on the right track. Scientists often reason in a counterfactual way, e.g., the statement $p(E|H) = x$ is often interpreted as "if H were the case, then there would be a chance of x that E happens". Can this be made relevant to the Problem of Old Evidence? And if yes, how? Here are some suggestions.

- Conditional Probability Primitivism. Some philosophers propose that conditional probabilities should be seen as *primitive* rather than as derived from the familiar formula $p(E|H) = p(E \wedge H)/p(H)$. Hájek (2003, "What Conditional Probability Could Not Be") argues convincingly that this definition leads into severe trouble, e.g., when the proposition upon which we condition has probability zero. This open happens in science, e.g., when conditioning on a particular parameter value in a statistical model. We should, or so the argument goes, discard the standard axioms of probability in favor of an axiomatization (e.g., Popper-Renyi) where conditional probability is taken as primitive. We could then say that the relevant likelihoods $p(E|H)$, $p(E|\neg H)$, etc. need not be equal to one if E is already known, and E could provide positive support for H (e.g., according to the Bayes factor). However, such a move seems to be *ad hoc* unless we provide a general story how conditional probabilities relate to suppositional reasoning and scientific reasoning in general.

- Scientific Reasoning as Suppositional Reasoning. The point of scientific confirmation judgments is not to determine our rational degree of beliefs that a hypothesis of interest H is true—after all, most scientific models are highly idealized and unlikely candidates for full truth. Rather, such judgments compare the explanatory power of scientific models for a given set of data and phenomena. For making such judgments, subjective assessments of the plausibility of different models are indispensable (e.g., in weighing the alternatives to a specific model). So all assessments of scientific evidence seem to be of a counterfactual nature, asking us to take an overarching model for granted. What scientists do in exploring the consequences of a scientific model resembles the Ramsey test for conditionals: would some particular observations follow from the model if we hypothetically assumed it as true? Such implications could also be probabilistic and take the form $p(E|H) = x$. Our thesis is therefore that conditional probabilities in scientific reasoning can be understood as the probabilities of conditionals. We investigate whether such an account can be consistently defended against a number of obvious objections.[1] On the positive side, this account aligns well with the view that conditional probabilities (rather than unconditional probabilities) should be taken as primitive. More generally, we aim at opening up perspectives for a theory of scientific evidence, and scientific reasoning more generally, as a theory of *suppositional reasoning*.

The paper rethinks Bayesian Confirmation Theory by exploring how far we can get following this road. Apart from answering the Problem of Old Evidence, we hope that it also helps us to address a frequently heard criticism of Bayesianism applied to scientific reasoning: that the degree of belief interpretation makes no sense in contexts where we are almost sure that all of our models are wrong.

[1]Consider, for instance, the multitude of triviality results that are supposed to prove the inadequacy of Stalnaker's Thesis $p(H \rightarrow E) = p(E|H)$.

Submission for 7th Progic Conference

Statistics and full belief

Jan-Willem Romeijn (Groningen)
Joint work with Hannes Leitgeb (LMU Munich)

Short abstract
This paper is concerned with the translation of statistical results into qualitative beliefs. Our approach to this problem relies on recent work concerning the relation between probabilistic and qualitative belief, based on a notion of stability or robustness. The stability view of full belief brings numerous attractive consequences for the logic of statistical claims, and it provides a natural context for appreciating the dependence of statistical methods on prior convictions, sampling plans, and the like. In particular, we consider classical hypothesis testing methods, and show how certain qualitative commitments, when combined with the stability view, entail probabilistic constraints that cannot always be satisfied. Conversely, it will become apparent that the conversion from probabilistic to categorical beliefs can be conceptually costly: it may involve violations of the likelihood principle and the principle of total evidence.

Extended abstract
Statistical science is concerned with probabilistic expressions of expectations, uncertainty and risk. But inevitably, there will be the occasional call from a decision or policy maker who asks what she can accept as facts, so as to choose her actions or back up her announcements. Obviously, one possible response will be to provide the caller with the full detail of the situation, laying out the probability assignment and eliciting the utility function of the decision maker, in order to give her quantified advice. However, it is well imagineable that such advice will be met with a request for clarity and simplicity. What can we safely assume at the moment? What claims are warranted?

This paper is concerned with the question what a statistician can offer in response to such a request. It concerns, in brief, the translation of statistical results into qualitative full beliefs. Our approach to this problem, which of course has a long history in the philosophy of statistics (e.g., Neyman, Dempster, Kyburg, and numerous others have contributed to the debate over this), relies on recent work concerning the relation between probabilistic and qualitative belief, based on a notion of stability or robustness: we only commit to claims that maintain a probability higher than a certain threshold after conditionalization upon all members of a relevant set of propositions (cf. Leitgeb 2013 and 2014). We consider a number of statistical methods and show how each of these can be connected to specific advice on what claims to commit to, and so include in our full belief state. Our reliance on the stability view of full belief brings numerous attractive consequences for the logic of statistical claims, and it provides a

natural context for appreciating the dependence of statistical methods on prior convictions, sampling plans, and the like.

For starters, a straightforward connection can be forged between Bayesian statistics and qualitative belief. The outcome of a Bayesian analysis is already a probability assignment over statistical hypotheses, so the afore-mentioned machinery of rational acceptance can be applied directly. Of course, it is a matter of debate if and how we can report full belief states after having done a Bayesian analysis. One issue is that we need to specify the relevant set of propositions, relative to which we call our belief in some hypothesis stable. Apart from that, some have advanced principled objections to the conversion to a full belief state. Radical probabilitists argue that this conversion presents an unwarranted and unnecessary impoverishment of the posterior probability assignment that a Bayesian analysis delivers. Against this, we argue that the practice of decision and policy may necessitate a move to qualitative judgements. Paraphrasing Dubins and Savage, we aim to explicate "what to commit to, if you must".

After looking at Bayesian methods we consider classical statistical methods, in particular Neyman-Pearson hypothesis testing (NPHT). We start by observing that classical methods do not employ probability assignments over statistical hypotheses, so that in first instance the criterion of stability cannot be applied. Our account relies on two further premises. First, we argue that even classical statisticians should feel comfortable with some qualitative commitments regarding statistical hypotheses, and we show that these qualitative commitments translate to specific constraints on a probability assignment over statistical hypotheses, i.e., specific interval-valued assignments. Second, we show that the procedure of NPHT is closely aligned to a Bayesian inference using interval-valued priors and posteriors, taking our cue from Romeijn (2010). With these two premises in place, we show that NPHT may warrant particular qualitative claims concerning statistical hypotheses. Depending on the set of propositions relative to which our belief in the hypothesis is required to be stable, we find non-trivial bounds on the priors and the thresholds.

When fitting NPHT into the framework of a stability-based conception of belief, it becomes apparent that the qualitative commitments warranted by the statistical procedure are highly context-sensitive. We end the paper by suggesting that this context-sensitive aspect of our account is precisely what we should expect to play a role, both from a general philosophical point of view, and more practically, when we are on the phone with the decision maker: we state what we robustly believe, but this is essentially relative to a context within which we are asked for our opinion.

Probabilistic logics with independence and probabilistic support

Dragan Doder, University of Luxembourg
Zoran Ognjanović, Mathematical Institute of SASA

Abstract. The main goal of this work is to present the proof-theoretical and model-theoretical approaches to probabilistic logics which allow reasoning about independence and probabilistic support. We extend the existing formalisms [7] to obtain several variants of probabilistic logics by adding qualitative operators for independence and support to the syntax. We axiomatize these logics, provide corresponding semantics, prove that the axiomatizations are sound and strongly complete, and discuss decidability issues.

Introduction. Independence is one of the main notions in probability theory. Two events, A and B are said to be independent w.r.t. a probability measure μ, if $\mu(A \cap B) = \mu(A)\mu(B)$ (or, more intuitively, $\mu(A|B) = \mu(A)$). Surprisingly, in spite of extensive development of various probabilistic logics in past decades, the notion of independence received little attention from the logical side.

Strongly related notion is probabilistic support [1]. We say that B probabilistically supports (or confirms) A, if $\mu(A|B) > \mu(A)$, i.e. $\mu(A \cap B) > \mu(A)\mu(B)$. There are several nonequivalent proposals for the measure of confirmation $c(A, B)$, all of them agreeing in a qualitative way: $c(A, B) > 0$ iff $\mu(A|B) > \mu(A)$ [4]. The notion of probabilistic support has drown particular attention for the famous discussion on its nature, after Popper and Miller [9] claimed that probabilistic support is deductive, not inductive (see e.g. [6]).

In this paper, we propose sound, complete and decidable logics for reasoning about independence and support. We start with probabilistic logics that extend classical propositional calculus with modal-like unary operators of the form $P_{\geqslant r}(\alpha)$, with the intended meaning "the probability of α is at least r" [7]. We extend the language with an additional binary operator $\alpha \nearrow \beta$ with the intended meaning "α confirms β or α and β are independent" (a weak form of support) as primitive operator. We introduce three formal systems: the logic LPP_2^{ind} without nesting of probabilistic operators, LPP_1^{ind} where nesting is allowed, and $LPP_2^{FR(n),ind}$ with a fixed finite range for probabilistic measure.

Related work. There are not many probabilistic logics in which independence and support are expressible. The reason for that lies in difficulties in combining high expressivity of the language and inference mechanisms. Fagin, Halpern and Megiddo [3] axiomatized logic with linear weight formulas (LWFs), i.e. Boolean combinations of the expressions of the form $r_1 P(\alpha_1) + \cdots + r_n P(\alpha_n) \geq r_{n+1}$, where r_i are rational numbers and α_i are propositional formulas. In the language, one can express the statement "conditional probability of α given β is at least r", but not linear combinations of conditional probabilities. There are two axiomatized logics which extend LWFs, that can capture independency. In [5], LWFs are enriched with operator that can express independence of sets of propositional letters. In [2], an axiomatization of the language that contains linear combinations of conditional probabilities is presented. Consequently, both independence and support are expressible.

The language with polynomial weight formulas (PWFs), which extends all previously mentioned languages, is also considered in [3]. A complete axiomatization of PWFs is presented in [8].

This work offers first axiomatization of the logics without arithmetical operations built into syntax, that can represent independence and support. Also, to the best of our knowledge, none of the previous attempts deals with nesting of probabilistic operators. On the other hand, we offer formalism in which we can express independence and support of both probabilistic and classical formulas.

Formalization. First we introduce LPP_2^{ind}. We extend the language of LPP_2 [7] by adding the binary operator $\nearrow$. Let For_C denote the set of all propositional formulas over the set of propositional letters $\mathcal{P}$. The basic probabilistic formulas have the following two forms: $P_{\geqslant r}\alpha$ and $\alpha \nearrow \beta$, where $\alpha, \beta \in For_C$ and $r \in [0, 1] \cup \mathbb{Q}$. The set For_P of all probabilistic formulas is the set of all Boolean combinations of the basic probabilistic formulas. Mixing of pure propositional formulas and probability formulas is not allowed. We introduce independency and support operators as abbreviations: $\alpha \perp \beta \equiv \alpha \nearrow \beta \wedge \neg\alpha \nearrow \beta$, $\alpha \uparrow \beta \equiv \alpha \nearrow \beta \wedge \neg(\alpha \perp \beta)$.

An LPP_2^{ind}-*model* is a structure $\langle W, H, \mu, v\rangle$ where W is a nonempty set the elements of which will be called worlds, H is an algebra of subsets of W, $\mu : H \longrightarrow [0,1]$ is a finitely additive measure, and v is a valuation which associated with every world $w \in W$ a truth assignment $v(w)$ on propositional formulas. We consider only so called measurable models, i.e. models such that each set $[\alpha] = \{w \in W \mid v(w)(\alpha) = \top\}$ is measurable. The satisfiability relation for $\alpha \in For_C$ is defined by: $M \models \alpha$ iff $v(w)(\alpha) = \top$ for all $w \in W$. For probabilistic formulas, we have

- $M \models \alpha \nearrow \beta$ iff $\mu([\alpha \wedge \beta]) \geqslant \mu([\alpha])\mu([\beta])$, $\quad M \models P_{\geqslant r}(\alpha)$ iff $\mu([\alpha]) \geqslant r$,
- $M \models \neg\alpha$ iff $M \not\models \alpha$, and $M \models \alpha \wedge \beta$ iff $M \models \alpha$ and $M \models \beta$.

The axiomatization of LPP_2 [7] contains an infinitary rule of inference (the key issue for probabilistic logics is in non-compactness, and, consequently, any finitary axiomatization is not strongly complete). We extend the axiomatization of LPP_2 with new axiom **Ax** $(P_{\geqslant r}\alpha \wedge P_{\geqslant s}\beta \wedge \alpha \nearrow \beta) \rightarrow P_{\geqslant rs}(\alpha \wedge \beta)$, and new infinitary rule

Ru From the set $\{\gamma \rightarrow ((P_{\geqslant r}\alpha \wedge P_{\geqslant s}\beta) \rightarrow P_{\geqslant rs}(\alpha \wedge \beta)) \mid r, s \in [0,1] \cup \mathbb{Q}\}$ infer $\gamma \rightarrow (\alpha \nearrow \beta)$.

We prove that the axiomatization is strongly complete with respect to the class of measurable models.

Then we will present the logic LPP_1^{ind} which extends LPP_2^{ind} so that iterations of the probabilistic operators are allowed. In that way we can express independence and support statements mixing classical and probabilistic formulas. For example, $\alpha \uparrow P_{\geq r}\beta$ is a formula of LPP_1^{ind}. Models of LPP_1^{ind} are structures of the form $M = \langle W, Prob, v\rangle$ where $Prob$ assigns to every $w \in W$ a probability space $Prob(w) = \langle W(w), H(w), \mu(w)\rangle$. Here $W(w)$ is a non empty subset of W, $H(w)$ is an algebra of subsets of $W(w)$ and $\mu(w) : H(w) \rightarrow [0,1]$ is a probability measure. In LPP_1^{ind}, we define $\models$ between worlds and formulas. For example, if $p \in \mathcal{P}$, then $M, w \models p$ iff $v(w)(p) = \top$, while $M, w \models \alpha \nearrow \beta$ iff $\mu(w)([\alpha \wedge \beta]_w) \geqslant \mu(w)([\alpha]_w)\mu([\beta]_w)$, where $[\alpha]_w = \{u \in W(w) \mid M, u \models \alpha\}$. We prove that the axiomatic system for LPP_2^{ind} is also sound and strongly complete for the logic LPP_1^{ind}.

Finally, we restrict LPP_2^{ind} to the logic $LPP_2^{FR(n),ind}$, focusing only to the probabilistic models have a finite fixed range $\{0, \frac{1}{n}, \ldots, \frac{n-1}{n}, 1\}$. We propose finitary axiomatization for the logic.

We also prove that all three logics proposed here are decidable.

References

[1] R.Carnap. Logical foundations of probability. The University of Chicago Press (1950).

[2] D. Doder, B. Marinković, P. Maksimović and A. Perović. A logic with conditional probability operators. *Publications de L'Institute Mathematique*, 87(101), 85–96 (2010).

[3] R. Fagin, J. Halpern and N. Megiddo. A logic for reasoning about probabilities. *Information and Computation* 87(1-2), 78 – 128 (1990).

[4] B. Fitelson. The plurality of Bayesian measures of confirmation and the problem of measure sensitivity. *Philosophy of science* 66, S362 – S378 (1999).

[5] M. Ivanovska, M. Giese. Probabilistic Logic with Conditional Independence Formulae. *Proc. of ECAI 2010*, 983–984.

[6] A. Mura. When probabilistic support is inductive. *Philosophy of science* 57 , 278–289 (2000).

[7] Z. Ognjanović and M. Rašković. Some first-order probability logics. *Theoretical Computer Science* 247 (1–2) , 191–212 (1990).

[8] A. Perović, Z. Ognjanović, M. Rašković, Z. Marković. A probabilistic logic with polynomial weight formulas. *Proc. of FoIKS 2008*, 239–252.

[9] K. Popper, D. Miller. A proof of the impossibility of inductive probability. *Nature* 302, 687–688 (1983).

The problem of iterated revision (Abstract, E.R.)

If we accept the idea that the problem of induction is a *dynamic* problem (vs. static), then it reduces to the problem of revising epistemic states. From a probabilistic perspective, there are two different paths: either adopting *objective Bayesianism* in the form of maximising entropy on the cumulated information (MaxEnt) P_I, or adopting (generalised) *orthodox Bayesianism*, in the form of minimising the Kullback-Leibler divergence on the new information with respect to the previous probability (MinKL) P_I^+. Formally, MaxEnt is MinKL starting from the equivocator $P_I = (P_=)_I^+$. However, both are competing epistemologies. The first holds that revision consists in re-evaluating the degrees of belief in the light of revised information. The second holds that revision consists in changing degrees of belief conservatively in the light of the newly accepted information. Both coincide if subsequent information grow stronger, i.e. $I_{i+1} \to I_i$. Yet, they diverge in the more general case, generating different liberalised induction schemes, where evidence may (but must not) be probabilistic. This paper addresses a problem encountered by both epistemologies - the *problem of iterated revision*. This is related to the question of how inductive support changes along the learning process. For the sake of simplicity, I leave the issue of convexity and closure aside and assume accepted information I (eg. linear constraints) to determine a closed-convex set of acceptable probabilities $\mathbb{P}(I)$ (if not, the convex hull is taken $\langle \mathbb{P}(I) \rangle$).

MinKL is uniquely defined if (the closed convex) $\mathbb{P}(I)$ (with I assumed consistent) admits at least one probability absolutely continuous to P. Therefore, MinKL is iterable if revision is restricted to the realm of regularity (the probability is regular and $\mathbb{P}(I)$ admits one regular probability). Iterability is interrupted, if the agent accepts at some point evidence which introduces some zeros, because all probabilities accepted by a future evidence which requires to raise one of these zeros will have infinite divergence to the former resulting probability. For conditionalisation, this is the problem of updating by a probability 0 element. The general problem (the impossibility to raise probability 0) is a form of the *expansion problem* (the impossibility to accept something which was previously rejected).

There are two options. One was suggested by Gärdenfors (1988) for conditionalisation: revise a probability by A by conditionalising the closest A-conditionalisable probability to P in an ordinal family (taking Popper measures is almost equivalent). This option encounters the *iteration problem*, because it violates the *principle of categorical matching* (Gärdenfors and al., 1995): the output of a revision should be of the same type as the input. The generated revision of a probability is a probability, not an ordinal family. A *rich* solution accepts rich epistemic states: replace ordinal families by ranked families and revise such families (Spohn, 1988, 2012). The *slender* solution sticks to probabilities. My slender solution (*ϵ-revision*) recommands to revise only after having taken a convex combination $(1-\epsilon)P + \epsilon P'_=$ of the previous probability P and the equivocator $P'_=$ on the newly introduced possibilitie. $1-\epsilon$ may be interpreted as the degree of conservatism or reluctance of the agent to introduce new possibilities. I sketch different options for a co-evolution of degrees of belief and the degree of conservatism.

A similar problem haunts objective Bayesianism. MaxEnt is uniquely defined as long as the cumulated information determines a (closed-convex) non-void set. This encounters the expansion problem, since the conjunction of two incompatible pieces of information determines the empty set. Two options are possible.

First, information may be cumulated by adopting a belief revision scheme: take the previous set of accepted probabilities as core and revise by the newly accepted probabilities. This solves the expansion problem, but only at the expense of introducing richer epistemic states (a well-ordered partition of the probability hypersurface) and faces ultimately the iteration problem. Revised information is a set of probabilities, not a well-ordered partition. I only see a rich option here: introduce a ranking function on the probability hypersurface and revise according to rank-revision.

Second, in the case of inconsistency, one may switch from a *conjunctive mode* of combining old and new information $\mathbb{P}(I) \cap \mathbb{P}(J) = \mathbb{P}(I \wedge J)$, to a *disjunctive mode*, and then take the convex hull $\langle \mathbb{P}(I) \cup \mathbb{P}(J) \rangle$ (Williamson, 2010, for some examples). This raises a new problem: the more mutually inconsistent evidence, the more equivocal the resulting degrees of belief. This seems plausible if evidence comes in a package, but implausible, if the evidence is ordered temporally or by reliability. The last (or most reliable) evidence should dominate here. The disjunctive mode also treats all pieces of evidence (old or new, reliable or less reliable) in the same manner. The ϵ-revision scheme enables to distinguish the weight (or reliability) of evidence: mix P_{I_1} with P_{I_2} and update by MinKL in $\langle \mathbb{P}(I_1) \cup \mathbb{P}(I_2) \rangle$ (a slight generalisation enables to distinguish differences in reliability).

Both approaches shed new light on the *progic* (Haenni et al., 2011) based on MaxEnt (Williamson, 2013). In Williamson's approach, $I \models J$[1] iff $\mathbb{P}^\delta(I^+, P_=) \subseteq \mathbb{P}(J)$ where $\mathbb{P}^\delta(I^+, P_=)$ are the δ-sufficiently equivocal probabilities of $\langle \mathbb{P}(I^+) \rangle$ (I^+ a chosen maximal consistent subset of I) with respect to $P_=$. The inference relation *implicitly* depends on the equivocator $P_=$ (or equivalently P_I), a bound of accepted upper divergence δ and an order allowing to choose a maximal subset, i.e. it depends on an epistemic state. This can be liberalised to any probability (replacing MaxEnt by MinKL): $I \models_{\delta,P} J$ iff $\mathbb{P}^\delta(I, P) \subseteq \mathbb{P}(J)$. Thereby the inference relation obtains a *dynamical dimension*: $I \models_{\delta,P}$ iff $I \models_{\delta,P_I^+}$. Additionally, the order of accepting information may be important, which leads to distinguish an ordered sequence ";" from a chunk ",": $I_1; I_2 \models_P$ (defined as $I_2 \models_{P_{I_1}^+}$) may differ from $I_2; I_1 \models_P$, because generally $(P_I^+)_J^+ \neq (P_J^+)_I^+$. Along the learning process, the reference probability P becomes subsequently replaced by updates and previously accepted information gets overwritten by new information (i.e. $I_1; \ldots; I_n \models_P$ is equivalent to $I_n \models_{(P_{I_1}^+ \cdots)_{I_n}^+}$). Replacing MinKL by ϵ-revision, the evolution of the inference relation also depends on the evolution of the degree of conservatism. I will present some dynamical laws for the evolution of these inference relation, based on transformational properties of MinKL (or ϵ-revision) and I will formulate a corresponding version for the disjunctive mode of MaxEnt. Following the outlined approachs, the inference relation evolves dynamically and depends *explicitly* on the actual epistemic state.

[1] $I := \{\phi_1^{X_1}, \ldots, \phi_n^{X_n}\}$ and $J = \psi^Y$, with ψ^Y meaning $Q(\psi) \in Y$.

References

Gärdenfors, P. (1988): *Knowledge in Flux. Modeling the Dynamics of Epistemic States*, Cambridge, MA: MIT Press.

— and Rott, H.(1995): "Belief Revision", in *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 4, Epistemic and Temporal Reasoning* (eds. D. M. Gabbay, C. J. Hogger and J. A. Robinson), Oxford: Oxford University Press, 35-132.

Haenni, R.; Romeijn, J. W.; Wheeler, G. and Williamson, J. (2011): *Probabilistic logics and probabilistic networks*, Dordrecht: Synthese Library, Springer.

Spohn, W. (1988): "Ordinal Conditional Functions. A Dynamic Theory of Epistemic States", in *Causation in Decision, Belief Change, and Statistics, Vol. 2* (eds. W. L. Harper and B. Skyrms), Dordrecht: Kluwer, 105-134.

— (2012): *The Laws of Belief: Ranking Theory and its Philosophical Applications*, Oxford: Oxford University Press.

Williamson, J. (2010): *In defence of objective Bayesianism*, New York: Oxford University Press.

— (2013): "From Bayesian Epistemology to Inductive Logic", project description.

Extended Abstract:

Belief Revision via Maximization of Expected Epistemic Utililty

---

In this paper, we provide a diachronic coherence constraint on (qualitative) beliefs and credences and discuss contrastive cases with the competing accounts. Our aim is to develop and explore a naïve, veritistic account of how *epistemically rational* agents should update their beliefs upon receipt of new information. We develop a substantiated and philosophically well-motivated version of what Foley [7] called the Lockean thesis – *i.e.* it is epistemically rational for an agent to believe φ iff she has credence in φ above some threshold. According to our constraint, an agent with a qualitative belief state $\mathbf{B}$ and a credal function b who then learns (with certainty) some proposition E is diachronically coherent iff her qualitative belief state after the update maximizes *expected epistemic utility* (EEU) relative to $b(\cdot \mid E)$.

Our scoring rule follows the one found in [5] in which an agent receives some positive epistemic utility $\mathfrak{r}$ for believing (disbelieving) a true (false) proposition and some negative epistemic utility $-\mathfrak{w}$ for believing (disbelieving) a false (true) proposition. This approach may be broadly motivated on Jamesian grounds [10], as it may be plausibly suggested that belief aims at the truth and disbelief aims at falsehood. Thus, an agent does well epistemically by believing truths and disbelieving falsehoods. However, an agent receives no utility or disutility form suspending on any proposition. While belief aims at the truth and disbelief aims at falsehood, suspension may be seen as opting out. For our current purposes, we leave the values of $\mathfrak{r}$ and $\mathfrak{w}$ as free parameters with only the restriction that $0.5 \leq \mathfrak{r} < \mathfrak{w}$.[1] In the present framework, this is motivated by the suggestion that it is *sometimes* epistemically permissible for an agent to suspend judgment on a proposition.

We assume that each rational agent has a probabilistic credence function, $b(\cdot \mid E)$. As has been argued elsewhere in the literature, this assumption can be independently motivated on veritistic grounds.[2] Thus, we calculate the EEU of an agents doxastic state based the scoring rule described above and the sum of the weighted sums of each of the propositions in some finite agenda, $\mathcal{A}$, where the belief state, $\mathbf{B}$, contains all and only the agents attitudes towards each proposition in the agenda:

$$EEU(\mathbf{B}, b) := \sum_{\varphi \in \mathcal{A}} \sum_{w \in W} b(w) \cdot u(\mathbf{B}(\varphi), w)$$

Based on these considerations, call an agent synchronically coherent iff she maximizes EEU. We may then follow [6] and say that an agent is synchronically coherent iff her doxastic state satisfies the following Lockean thresholds for an EEU maximizing agent:

Theorem 1 *An agent with credal function b and beliefs ($\mathcal{B}$), disbeliefs ($\mathcal{D}$), and suspentions ($\mathcal{S}$) will maximize EEU relative to her credence function iff for every $\varphi \in \mathcal{A}$*

$$\begin{array}{lll} \mathcal{B}(\varphi) & \textit{iff} & b(\varphi) > \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}} \\ \mathcal{D}(\varphi) & \textit{iff} & b(\varphi) < \frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}} \\ \mathcal{S}(\varphi) & \textit{iff} & b(\varphi) \in \left[\frac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}}, \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}\right] \end{array}$$

[1] Although we do not presently place any further restrictions on the values of these parameters, this is not because there is nothing substantive to say here. William James suggested that the relative values of these parameters depends merely on the "dispassional nature" of the agent; however, Pruss [15] argues that $\mathfrak{r}$ should be at least 2.588 times the value of $\mathfrak{w}$ based on some plausible assumptions about the nature of the scoring rule. Such discussions, while important, are incidental for our current project and so we remain neutral on this issue for ecumenical reasons.

[2] For a survey of these arguments, see [14].

Subsequently, we rely on Bayesian conditionalization (for which further veritistic arguments have also been given [9]) to constrain the updates of full beliefs. An agent with an initial belief set, **B**, and a belief set after learning E (with certainty), $\mathbf{B}'$, is diachronically coherent iff $\mathbf{B}'$ maximizes EEU relative to $b(\cdot \mid E)$. Thus, we may define the qualitative belief revision $\mathbf{B}' = \mathbf{B} \circledast E$ as follows:

$$\mathbf{B} \circledast E := \{\mathsf{B}\varphi, \mathsf{D}\psi, \mathsf{S}\theta : b(\varphi \mid E) > \tfrac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}, b(\psi \mid E) < \tfrac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}}, b(\theta \mid E) \in \left[\tfrac{\mathfrak{r}}{\mathfrak{r}+\mathfrak{w}}, \tfrac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}\right]\}$$

This sort of qualitative updating will have a number of features that differ from standard qualitative updates.

Of particular interest, we note that $\circledast$ will not satisfy the following principle found in [12]:

(P2) If the agent already believes X, then updating on the piece of evidence X does not change her system of (all-or-nothing) beliefs at all.

This principle is satisfied by the two most prominent competitors for qualitative updates, AGM belief revision [1] (and derivatively Leitgeb's Humean thesis) and Lin and Kelly tracking [13]. Not only do we provide some reason to resist **(P2)**, but we derive a general lower bound on the value of $b(\psi \mid \varphi)$ where $b(\varphi) > \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}$, $b(\psi) > \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}$.

Theorem 2 *Where* $b(\varphi) > \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}$, $b(\psi) > \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}}$, *and* $0.5 \leq \frac{\mathfrak{w}}{\mathfrak{r}+\mathfrak{w}} < 1$, *the lower bound on* $b(\psi|\varphi)$ *is* $\frac{\mathfrak{w}-\mathfrak{r}}{\mathfrak{w}}$.

Interestingly, this means that when the when the threshold is 0.5, the lower bound is 0, while as the threshold approaches 1, the lower bound approaches the threshold. However, **(P2)** will be satisfied just in case the agent has extremal credences in the relevant propositions. Perhaps this helps explain the intuition regarding why the principle is a good one.

The constraint that we develop in this paper is admittedly a naïve first pass at providing a veritistic diachronic coherence constraint on beliefs and credences. There are a number of important ways that the account will need to be developed in order to accommodate more cases. For example, we do not intend to accommodate dependencies between truth-values of φ, $\mathsf{B}\varphi$, and b as in the cases found in the recent literature[3] Moreover, we only develop a diachronic constraint for when agents acquire information with certainty. It remains to extend this treatment to accommodate Jeffrey conditionalization. Finally, we have yet to establish a purely qualitative representation of the revision operation. There is, however, hope that this will be possible. Recently, van Eijck and Renne [16] built a logic for Lockean beliefs where the threshold is 0.5. Our aim in this paper is simply to lay out the minimal naïve account of qualitative belief update via maximization of expected epistemic utility with the hope that this will serve to spur further investigation of this approach.

References

[1] C. E. Alchourron, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2):510–530, June 1985.

[2] M. Caie. Belief and indeterminacy. *Philosophical Review*, 121(1):1–54, January 2012.

[3] M. Caie. Rational probabilistic incoherence. *Philosophical Review*, 122(4):527–575, 2013.

[4] J. Carr. Epistemic utility and the aim of belief. 2014.

[5] K. Easwaran. Dr. Truthlove, or how I learned to stop worrying and love Bayesian probabilities. http://dl.dropbox.com/u/10561191/Unpublished/Truthlove.pdf, 2014.

[6] B. Fitelson. *Coherence*. Oxford University Press, Oxford, Forthcoming.

[7] R. Foley. The epistemology of belief and the epistemology of degrees of belief. *American Philosophical Quarterly*, 29(2):111–121, 1992.

[8] H. Greaves. Epistemic decision theory. *Mind*, forthcoming.

[3]For example, as discussed in [4], [2, 3], [8], and [11].

[9] H. Greaves and D. Wallace. Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, 115(459):607–632, July 2006.

[10] W. James. The will to believe. *The New World*, 5:327–347, 1896.

[11] J. Konek and B. Levinstein. The foundations of epistemic decision theory. 2014.

[12] H. Leitgeb. Reducing belief *Simpliciter* to degrees of belief. *Annals of Pure and Applied Logic*, 164:1338–1389, 2013.

[13] H. Lin and K. Kelly. Propositional reasoning that tracks probabilistic reasoning. *Journal of Philosophical Logic*, 41(6):957–981, 2012.

[14] R. Pettigrew. Epistemic utility arguments for probabilism. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Winter 2011 edition, 2011.

[15] A. R. Pruss. Independent tests and the log-likelihood-ratio measure of confirmation. *Thought*, 3:124–135, 2014.

[16] J. van Eijck and B. Renne. Belief as willingness to bet. 2014.