

BELIEF AND CONTEXTUAL ACCEPTANCE*

Second version – June 16, 2008

Eleonora Cresto

CONICET (National Council for Scientific and Technical Research), Argentina

eleonora.cresto@gmail.com

ABSTRACT

In this paper I develop a strategy for representing epistemic states and epistemic changes that seeks to be sensitive to the difference between voluntary and involuntary aspects of our epistemic life, as well as to the role of pragmatic factors in epistemology. The model relies on a particular understanding of the distinction between *full belief* and *acceptance*, which makes room for the idea that our reasoning on both practical and theoretical matters typically proceeds in a contextual way. Within this framework, I discuss how agents can rationally shift their credal probability functions so as to consciously modify some of their contextual acceptances; the present account also allows us to represent how the very set of contexts evolves. Voluntary credal shifts, in turn, might provoke changes in the agent's beliefs, but I show that this is actually a side effect of performing multiple adjustments in the total lot of the agent's acceptance sets. In this way we obtain a model that preserves many pre-theoretical intuitions about what counts as adequate rationality constraints on our actual practices – and hence about what counts as an adequate, normative epistemological perspective.

1. Introduction. The belief/ acceptance distinction

In this paper I propose a strategy for modeling the epistemic state and epistemic changes of a particular agent at a given time. The model seeks to

- (a) illuminate the extent to which there is room for pragmatic factors in epistemology;

* Part of the material included in this paper benefited from conversations with John Collins, Isaac Levi and Achille Varzi, all of whom provided useful criticism and advice for improvement. Early drafts have been presented at the *Vth Principia Symposium* (Florianópolis, August 2007), and at a discussion session at the *GAF (Grupo de Acción Filosófica)*, Buenos Aires, September 2007). I am grateful to the participants for very helpful comments and suggestions, especially to Anjan Chakravartty, Bas van Fraassen and Bradley Monton (in Florianópolis), and to Eduardo Barrio, Ramiro Caso, Javier Castro Albano, Milton Laufer, Eleonora Orlando, Federico Pailos, Gonzalo Rodríguez-Pereyra, Florencia Rimoldi, Laura Skerk, and Ezequiel Zerbudis (in Buenos Aires).

- (b) be sensitive to the difference between voluntary and involuntary aspects of our epistemic life; and
- (c) explore the extent to which our reasoning about both epistemic and practical matters proceeds in a contextual way.

I will argue that, insofar as these goals are fulfilled, we obtain a representation tool that preserves many pre-theoretical intuitions about what counts as adequate rationality constraints on our actual practices – and hence about what counts as an adequate, normative epistemological perspective. In addition, I hope to show that the particular account that I offer here exhibits several technical advantages (to be mentioned in due course) over alternative ways of proceeding. To carry out this project I shall rely on (some brand of) a cognitive decision theoretic framework, and I shall suggest a particular way of construing the distinction between *believing* a statement (idea, proposition, hypothesis or theory) and *accepting* it.

The notion of belief has been credited with incompatible features on different occasions; in particular, we need to reconcile somehow the Humean intuition that we cannot believe at will, with the equally strong intuition that agents routinely make non-deductive inferences and seek to change their minds on the basis of the conclusions of such inferences, while they also seek to convince others through rational conversation. We might be tempted to think that part of the problem here is that different authors have embraced very different conceptions of belief, *ontologically speaking*. For example, if beliefs are mainly characterized as *epistemic commitments* (as in Levi 1980, 1997, 2004 – to mention a few), involuntarism does not look too promising; the very idea of commitment embodies an irreducible normative element, and, *prima facie*, it seems to imply that we can be held responsible for the beliefs we have. By contrast, if beliefs are understood first and foremost as *dispositions* of some sort, involuntarism becomes more plausible: it seems that dispositions can well be acquired (and maintained) without our willing this to happen; in general, while it still makes some room for normativity, a conception of beliefs as dispositions exhibits a more naturalistic bent than one cashed out entirely in terms of commitments.¹

¹ Theories that conceive of beliefs in terms of commitments typically claim that commitments are not fully reducible to dispositions (cf. again Levi, 1997, ch.1); however, they may well entail that if an agent is, say, committed to *p*, then she has all sorts of dispositions to act or feel as if *p* – depending on the details of the

Still, taking a stance on the ontological debate does not suffice to settle the issue as to whether particular putative features should or should not hold. For instance, we might well be able to offer a theory of beliefs according to which we find ourselves, as it were, having or lacking certain specific commitments. In other words, the potential inconsistencies that we may find among alternative characterizations of beliefs are not guaranteed to go away once we clarify our ontological assumptions.

In the light of this, we might want to stipulate a belief-acceptance distinction in order to restore consistency and alleviate the tension. As is well known, different variants of the belief-acceptance distinction have been proposed during recent decades, with the aim of solving very different problems. Indeed, a quick look at the literature shows that there is no uniform way of understanding these concepts. Some authors, for instance, have emphasized that believing that something is the case usually entails being convinced of its truth, and have pointed out that at times we would like the connection with truth to be relaxed. Probably the best-known example of this perspective is found in Bas van Fraassen's discussion of the difference between being fully convinced of the truth of a given hypothesis and coming to accept it in order to keep on working along a particular research line; cf. (van Fraassen, 1989, 2002); similar motivations can be found in (Maher, 1993), although Maher's and van Fraassen's accounts do not yield extensionally equivalent pairs of concepts. In van Fraassen's case, in addition, the idea of acceptance is meant to help agents avoid committing themselves to the truth of hypotheses or theories that refer to unobservable entities. On the other hand, authors such as Jonathan Cohen (Cohen, 1992) have stressed that beliefs are not voluntary; as opposed to acceptances, they grow in us passively (similarly, cf. Lehrer, 2000). Still others, such as (Stalnaker, 1984) or (Bratman, 1992), have suggested that acceptances, as opposed to beliefs, refer to those propositions that we are only willing to assert in particular contexts.²

view. Similar qualifications may apply to comparisons between theories of belief as commitments/dispositions and descriptions of beliefs in terms of mental states. In any case, I should warn the reader that here I am not attempting to draw a complete map of different perspective on the nature of belief, and this is certainly not the place to dig into the difficulties and potential advantages of alternative positions. As I shall emphasize below, the model I build in this paper seeks to be neutral among different conceptions of what beliefs actually *are*.

² Cf. also the articles in Engel (2000). For yet other proposals see Kaplan (1996), Tuomela (2000), or Da Costa and French (2003).

Which way shall we go, then? In subsequent pages I shall develop a representation strategy to model beliefs and acceptances, which seeks to fulfill the goals mentioned at the beginning of this section.

2. Levels of analysis

Before presenting the structure of the model let me distinguish four possible levels of analysis; the distinction will help me clarify the scope of the paper. We find, at the very least:

- (A) The level of the agent's *real* epistemic state, which may be constituted by epistemic attitudes such as doubts, beliefs, and acceptances, among others – depending on our favorite theory.
- (B) The level of the agent's attitudes towards statements of her own language, such as acceptance, rejection, or suspension of judgment.³
- (C) The level of the agent's semantic assumptions about statements of her own language (in a way to be clarified soon).
- (D) The representation level.

The agent's real doubts, beliefs and acceptances, at level (A), are obviously compatible with more than one linguistic manifestation of (at least part of) such states, at level (B); in addition, elements at levels (A) and (B) are compatible with more than one modeling strategy. As for level (C), notice that by merely looking at the level of the agent's attitudes towards statements of her own language (level (B)) we do not know whether the agent is, or is not, committed to bivalence. If the agent is not committed to bivalence, the set of sentences she rejects [accepts] and the set of sentences she takes to be false [true] may not coincide, as she might accept or reject statements that she takes to be semantically undefined. Independently of this problem, an agent might accept [reject] a statement of her language and suspend judgment *on its truth-value*. Consequently, a representation

³ Notice that, depending on our favorite theory, this level might involve the identification of at least two different ways of “accepting” a sentence – in agreement with the identification of different epistemic attitudes in (A).

apparatus could give us all the information we deem relevant about level (B) without thereby providing enough information about what statements the agent takes to be true, false, or undefined (in case there are such undefined statements according to her favorite semantic theory).⁴

In this paper *I shall focus on level (D)*, with the aim of capturing some phenomena (in particular, with the aim of representing a number of distinctions) that are already intuitively clear at level (A). Let me emphasize, then, that I shall not attempt to develop a comprehensive, complete theory about the agent’s real epistemic state; in particular, I shall bypass what we might dub “the ontological question” on beliefs: the analytic tool I offer here will be compatible with many different approaches on the exact nature (at level (A)) of the potential epistemic states that are being so represented (such as sets of commitments, elements of a Boolean algebra, sets of dispositions, or neurological events, to mention a few – where these options need not be pairwise incompatible). Likewise, I shall not be particularly concerned with levels (B) or (C). As we shall see, I shall model epistemic attitudes by means of sentences of a representation language *L*, such that the sentences of *L* (at level (D)) *might* be taken to be idealizations (or perhaps suitable translations) of sentences of a language the agent speaks; nevertheless, an analysis of the agent’s attitudes about sentences of her own language, or of her semantic assumptions, will not be a goal in itself.

In the next section I shall argue that, according to well-entrenched pre-theoretical intuitions, there are pragmatic and contextual factors that shape crucial features of the agent’s epistemic life – at the level of her real epistemic state.

3. Contextual assumptions

Consider an agent who:

- (i) tries to convince others about the truth of a particular set of claims, or about the correctness, or the convenience, of a particular course of action; or

⁴ In Cresto (2008b) I develop a proposal that relies on the explicit distinction between the agent’s semantic assumptions and the semantics embedded in the representation tool.

- (ii) deliberates to figure out whether to perform a particular action; or
- (iii) deliberates to figure out whether to endorse a particular claim (regardless of whether she foresees the possibility of putting the result of the deliberation to some practical use)

It is clear that, in each case, the agent's reasoning is bound to take place within a particular cluster of circumstances. It is also clear, I think, that in each case the agent finds herself taking for granted a particular set of assumptions. In particular, she proceeds under the assumption that a number of ideas or propositions (I shall leave the terminology here intentionally vague) *are true*. She is *certain* about them, in the sense that:

- (a) as far as she is concerned, they are not open to discussion at that particular moment (i.e., at the time of proceeding with her reasoning);
- (b) they determine what is and what is not epistemically possible for her at the time of proceeding with her current problem; under a different terminology, they constitute her "standard for serious possibility" for the particular argumentative exercise with which she is involved at the moment: ideas or propositions that are logically compatible with her set of assumptions are judged possible; those that are not, are judged impossible.⁵

Typically, the agent also takes for granted the correctness of all sorts of moral and aesthetic judgments, as well as more encompassing views about rationality or agency, among other things. In what follows I shall not be concerned with them, but only with what we might call "*epistemic assumptions*."⁶ In addition, at the time of engaging in a particular reasoning agents are more or less *uncertain* about a number of ideas or propositions.

So far, the picture I have delineated here is not too controversial, I hope. Let me add now a further element. As I see it, not only does the agent's reasoning take place within a particular set of circumstances (such as the type of audience she is addressing, or her degree of emotional involvement with the topic, among many others) but also, those particular circumstances *shape* her reasoning, so to speak. By this I mean that certain specific features

⁵ The reader should be warned that this description proceeds from a first person perspective – at a phenomenological level, so to speak. It will be duly qualified in the next section.

⁶ By "epistemic assumptions" I mean to refer to elements (propositions, statements or ideas) that can be deemed true or false – as opposed to assumptions about what the agent takes to be right or wrong, or nice/interesting/ funny, etc. In particular, probability judgments are not epistemic assumptions; I shall come back to the last point below (thanks to Bernhard Nickel for pressing the issue).

of the agent's reasoning depend on (and can be seen to vary with) the type of circumstances. Part of the explanation for this phenomenon is that the peculiar circumstances within which the agent proceeds may affect a number of relevant variables (such as the degree of evidential support the agent deems to be sufficient in order to assent to a particular claim), but, more importantly, they may also affect *the set of epistemic assumptions that she is ready to take for granted* for the specific problem under consideration. In other words, I submit that part of the epistemic assumptions that an agent takes for granted at a particular time vary with the circumstances, where this variation does not constitute a change in view.

Let me give some examples to motivate this proposal.⁷ Suppose agent *X* reflects on whether illness *n* is the result of a single factor *f*. *X* can find it reasonable to accept this idea as true in order to go on studying the genesis of other illnesses; nevertheless, at the time of devising a public health policy she might think that it is necessary to be much more cautious, and perhaps she might conclude that it is better to suspend judgment and proceed as if the truth or falsity of the claim were still an open question. In this example, it is clear that there is a sense in which the agent is not *fully* convinced that *n* is caused by *f*. However, she is also determined to take the aforementioned causal relation for granted for an important part of her life – say, at the time of speaking at a symposium on the causes of illness *n*, discussing the topic with her colleagues after lunch, designing new experiments at the lab, or drawing some (non-deductive) conclusions on the likeness that *f* also causes *m*, among many other things. On the other hand, *X* takes for granted that, if untreated, illness *n* is mortal in more than 90% of the cases; as opposed to the previous example, at the moment she can think of no set of circumstances under which she would feel uncertain about this ratio: she is fully convinced that this is how things actually are.

For a different example, at the time of interacting with my neighbor I am certain that he is a student at the local college; I also take this idea for granted at the time of deciding which Christmas present to buy him – but I might find myself suspending judgment were I to be asked about this particular topic in court. This is not, however, to be

⁷ For a more elaborate example, see the Appendix.

thought of as an epistemic change. It is not that I have changed my mind: when I speak to him, I do not take seriously the possibility that he is not a college student.

Let the label “context” refer to the cluster of circumstances and assumptions (epistemic and non-epistemic) within which the agent’s reasoning and rational behavior takes place, and which makes it the particular instance of reasoning, or token of behavior, it actually is. My suggestion, then, is that the agent’s arguing, deliberating, deciding (on theoretical and practical matters), as well as her acting on the basis of prior deliberations and decisions, is contextual, and that contexts are partly defined by sets of epistemic assumptions – or sets of “certainties”.⁸ Thus, I shall say that agents find themselves holding specific epistemic assumptions *for specific contexts of action and deliberation* – or contextual certainties – in addition to epistemic assumptions for all contexts they deem relevant. Correspondingly, I shall say that at times agents consider the possibility of holding *new certainties* for particular contexts of action and deliberation, and not for others. For pragmatic reasons, as this paper is focused on the structure of *epistemic states*, I shall also say that if contexts *i* and *j* are constituted by the same sets of epistemic assumptions, then $i=j$.

Notice that epistemic assumptions, as I conceive of them, lead agents *to act* in certain definite ways. This suggests that the phenomenon I am trying to describe here is not to be confused with *hypothetical reasoning*.⁹ Moreover, when we say that an epistemic decision has been taken under hypothetical assumptions, we seem to imply that such a decision is somehow provisory. We seem to imply that at some point we should seek to “cancel” the assumptions, so to speak. But none of these connotations is adequate. Contextual epistemic decisions are not necessarily provisory; they are not to be thought of as decisions an agent takes temporarily until she makes up her mind as to which full convictions to hold. By contrast, I take it that *the existence of multiple decision contexts is an essential part of an agent’s typical epistemic state*.

In the light of the above, it should be clear that not every situation in which an agent acts *as if p* will be indicative of her being certain of *p* at the level of her *real* epistemic

⁸ For a related, but different, conception of contextual assumptions, see Nozick (1993).

⁹ In this I follow Bratman (1992), p. 9, against Stalnaker (1984).

state.¹⁰ Suppose I am trying to convince you of the truth of p ; in order not to beg any questions, I pretend to be in suspense about p during the course of our conversation. But this is just *pretense*; as a matter of fact, the context in which my argument proceeds is clearly a context in which I am certain that p : that's why I bother offering an argument for p in the first place. The same goes for belief-contravening assumptions for the sake of the argument – the type of assumptions one finds in *reductio* arguments. Again, the starting point of a *reductio* should not be taken to reflect our actual epistemic state. We can compare this phenomenon with a related one: suppose a dangerous criminal points at me with a gun and demands that I say “I am not sure whether I am a human being”. In that situation I will probably utter those words, just to avoid being killed. But that does not mean that I am not convinced of my being human. In short, at times it might be prudent for an agent not to disclose her real epistemic state, for whatever reason; as a special case of this general fact, at times it might be convenient to pretend (perhaps explicitly, as when I am arguing with a friend) that one has doubts and certainties one does not have.

In what follows I shall speak of “acceptances” to refer to the agent’s epistemic assumptions in some particular context, whereas I shall speak of “full beliefs” (or “beliefs”, for short) to refer to epistemic assumptions in all relevant contexts of action and deliberation, as far as the agent is concerned. Under this definition, a belief is a special type of acceptance. Notice also that within each context an agent will typically have more epistemic assumptions than full beliefs. Finally, we obtain that an agent may hold “contextual doubts” (doubts that are relative to one specific context but not others), as much as “absolute doubts.” I shall have much more to say about doubts in the next section.

I will not give further precisions on what contexts *are*; in particular, I do not think we can give a characterization of contexts that may in turn clarify the concept of contextual acceptance in a non-circular way, as any description of a context is bound to make reference to those things the agent takes for granted at the time of acting or reasoning. Related to this, a more substantive account of the concept of acceptance would force us to take a stance on the ontological problem of epistemic states, which is beyond the limits of this paper. In any case, I trust that, even if the notion has not been given a precise

¹⁰ Thank you to Bernhard Nickel for pointing this out to me.

definition, it is nonetheless pre-theoretically intuitive, and hence it will suffice for our present purposes.

Still, a number of clarifications are in order. First, what I have just referred to as a “context” is not to be understood in the way so-called contextualism in epistemology does.¹¹ According to the contextualist, contexts are individuated by reflecting on the knowledge ascriptions that come out true or false in each case. By contrast, here I shall not be concerned with the concept of *knowledge*, so I shall not be concerned with whether particular knowledge ascriptions (including possibly self-ascriptions) are true or not. More generally, even if we assume that belief (or perhaps *rational* belief) can be somehow equated with knowledge (an equation that some brands of pragmatist epistemology are prone to endorse),¹² different contexts in my sense are not to be thought of as involving higher order assessments of epistemic claims of a lower order. If anything, my understanding of contexts is closer to what Hawthorne (2004) or Stanley (2005) call “interests” (of the primary epistemic subject, rather than of those who assess particular epistemic assertions). Related to this, it should be noticed that in this paper I will not be especially concerned with iterated epistemic attitudes, or with the concomitant discussion about issues such as transparency or authority (for instance, can we assume that if agent *X* accepts that *p* in context *i*, then *X* believes – in all contexts – that *X* accepts that *p* in context *i*?) – although it seems safe to assume that the present account can be easily extended to take care of these topics.

Second, I have suggested that beliefs are to be understood as epistemic assumptions in all relevant contexts. We may need to say a few words on the notion of *relevance* at stake. Which contexts are relevant, exactly? Is there such a thing as a skeptical context, for instance – that is, a context in which we feel particularly cautious about everything? Or, if not, can it be the case that for every contingent potential assumption an agent can conceive of, there is a particular context in which it is absent? Even though in this paper I do not

¹¹ Cf. Cohen (2000), or DeRose (2002), among others.

¹² Many formal theories of belief revision (such as AGM theory; cf. Alchourrón *et. al.*, 1985; or Gärdenfors, 1988) equate knowledge with true full belief, as assessed from the agent’s perspective. I feel sympathetic to this position, although I will not defend it here. The rationale for this view can be found in contemporary elaborations of Peircean epistemology – the so-called “Belief-Doubt Model”. Within the Peircean framework, we need not justify prior beliefs in order to have knowledge; only changes are in need of justification; cf. for example Levi (1997), p. 4, Bilgrami (2000), pp. 251-257, or Bilgrami (2004).

intend to examine the concept of knowledge (which is the typical skeptical target)¹³ we might still fear that, if we are unable to attain a state of psychological certainty, then *strictu sensu* we cannot be said to hold full beliefs in the first place – justified or otherwise. If this were the case, it might be contended that acknowledging the legitimacy of a skeptical context, or of scattered “cautious” contexts for every contingent potential assumption, suffices to reduce bona fide beliefs to beliefs about logical truths.¹⁴

As I can see it, it is not our task, as theoreticians, to legislate which full beliefs agents should have (or lack), as a matter of rationality – nor to tell agents which contexts are legitimate.¹⁵ Which contexts are relevant for an agent at *t* depends in part on which sets of circumstances the agent can conceive of and deem possible to affect her reasoning – for some possible piece of theoretical or practical reasoning – at *t*. And there are no recipes that could tell us which sets of circumstances the agent should take into account at this point. In any case, I am also convinced that not all skeptical doubts are alike, and that no serious epistemological reflection should gloss over this fact – some such “doubts” are simply not compelling to most of us. Speaking by myself, I am ready to acknowledge that, at least in some contexts, I do not feel certain about the truth of claims concerning the occurrence of future events, so I will happily concede that I do not fully believe them, whereas I can identify no context of my life in which I have real doubts about my being human – as opposed to, say, a brain in a vat.

Third, let me address a final concern about the very idea of paying attention to contexts and contextual acceptances. It could be objected that, even if agents were in fact prone to taking different ideas for granted in different contexts, it is far from obvious that the distinction between belief and acceptance deserves serious consideration within a normative account. After all, the fact that real agents are often inconsistent is seldom taken to be a sign that we should develop an account that preserves this trait. The two cases, however, are not analogous. I *am* sympathetic to the thought that standards of rationality function as regulative ideals, and that they impose constraints on the range of acceptable

¹³ Notice, however, that if we take our model to refer to *justified* belief (rather than to belief *simpliciter*), we are again liable to skeptical threats.

¹⁴ The very set of logical truths could also be questioned, of course, but I shall not enter into this problem here.

¹⁵ As it is apparent, in this I follow a Peircean-style approach. Cf. footnote 13 above.

models; in particular, I agree that epistemic rationality imposes the ideal of consistency and deductive closeness on us, in the sense that we do long for epistemic states that exhibit such features, even if in practice they are unattainable. By contrast, we do not long for a unique standard for serious possibility at a given time t – at least not in the same sense in which we long for consistency. Yes, we do feel the pressure to make up our mind about particular topics in as much contexts as we could conceive of, so as to obtain, in the limit, no undecided questions left (in which case the distinction between different contexts as defined here would disappear) – but this is not analogous to the quest for consistency. Insofar as the idea of rationality and rational agency does not force us to say that we should have not doubts whatsoever, it does not force us to say that we should hold a unique set of certainties either. What we can conclude from here is that, in a normative model, consistency and deductive closeness should be relativized to contexts.

In the next section I shall present the basic structure of the model I favor, which seeks to capture the main intuitions delineated here.

4. Modeling beliefs and contextual acceptances

Let me suggest a model – at level (D) – to account for the epistemic distinctions discussed in section 3. The first question we should ask ourselves at this point is: should we take the informal description of section 3 at face value? Notice, in particular, that our discussion from section 3 seems to imply that full beliefs and acceptances are somehow more basic than *uncertainties*. Should we seek to preserve this consequence? To put it differently, should we take this feature of our previous description to reflect an important trait of the agent's real epistemic state? The answer, I take it, is 'no.' To many agents, beliefs and acceptances might indeed *seem* more basic –from a *phenomenological* point of view – but this does not entail that they are *ontologically* prior; in particular, I shall not assume that they carry ontological priority over doubts.¹⁶ The relevance of these qualifications will be apparent shortly.

¹⁶ I shall come back to a more detailed discussion of this point by the end of this section.

I shall begin by assuming that we can represent the potential certainties of a particular agent, at a particular time, as sentences of a suitably regimented language L . Agent X 's epistemic state at t will be modeled by a convex set $\Delta_{X,t}$ of probability functions P_k over the sentences of L . (I shall get rid of sub-indices when there is no risk of confusion). Now let \mathbf{T} be the set of all theories T_i of L , and define $f: \Delta \rightarrow \mathbf{T}$ such that $f(P_k) = \{\alpha \in L: P_k(\alpha) = 1\}$.¹⁷ In other words, we define a function from Δ to \mathbf{T} that maps each probability measure in Δ to the largest theory whose elements are assigned probability 1 by that measure. Some members of Δ will pick out the same element in \mathbf{T} – and some theories will not be selected at all, of course. Then, for any $T_i \in \mathbf{T}$, T_i is an acceptance set iff it is selected by some member of Δ . In turn, each context i can be correlated with a particular $\Delta_i \subseteq \Delta$ that contains all elements P_k which pick out T_i . By extension, we shall also say that Δ_i is correlated with $T_i = A_i$, or the set of acceptances of the agent in context i . It is easy to see that the Δ_j (the subsets of Δ correlated with particular contexts) constitute a partition, and that each Δ_j is convex as well.¹⁸

Next, define $K = \cap_j A_j$, for all j – hence, every function in Δ assigns probability 1 to every element in K . K represents the agent's full beliefs. Notice that K is bound to be itself an acceptance set: due to the convexity of Δ , there should exist a subset Δ_k with functions that pick out K ; let me call k the *minimal context*. More generally, there might be further acceptance sets (in addition to K) embedded in others, though this is not mandatory.

In addition, we can also define $M = \cup_j A_j$, for all j . M stands for the agent's total set of acceptances; all elements in M receive probability 1 by some (but not necessarily the same) function in Δ . Notice that I have not required that the A_i be pairwise consistent; in

¹⁷ Thanks to an anonymous referee for the *FEW 2008* for suggesting this formulation.

¹⁸ By demanding convexity we go some way towards answering a traditional criticism to standard Bayesianism – to wit, that it is unwise to assume that agents can be credited with precise probability assignments. In addition, convexity can be important at the time of solving standard decision theoretic problems. As is well known, when different probability measures yield incompatible options with maximum expected utility, intermediate probability values may enable additional options (say, a second best) to be eligible as well – which, depending on the details, may be seen as a natural way to commensurate alternative rankings. An even more compelling reason to demand convexity may be found in the thought that, when trying to reach a consensus between incompatible credal states, we need to be able to move to a position of suspense. I shall not consider this type of revision of a credal state here. On these points cf. paradigmatically Levi (1974, 1980).

any case, this fact does not enable agents to deduce explicit contradictions, because M need not be deductively closed – hence, the problematic statements (in case there are any) can remain isolated from one another, so to speak.¹⁹ Notice also that M will be correlated with a special context if and only if there is some probability function that picks it out, which need not exist (in any case, it is clear that M can only be correlated with a special context if it is consistent and deductively closed).

Within this setting, we have multiple ways in which we can say that an agent is in doubt. I shall say that sentences of L logically compatible with A_i but which do not belong to A_i represent *contextual doubts* regarding context i , whereas sentences of L that are logically compatible with every A_j but do not belong to M represent *full doubts* of the agent.²⁰ Clearly, contextual and full doubts are probabilistically ordered.

Let me stress that sentences of L_X (and hence of M_X) represent *basic* potential certainties of X rather than, say, judgments of epistemic possibility, or subjunctive and counterfactual conditionals. Conditional and modal claims can be assumed to be *licensed* by the structure of the agent's epistemic state, and hence of a derivative nature.²¹

Let me also emphasize that each P_j in Δ_i (for some context i) is *not* meant to describe a particular way in which things could be. Rather, each P_j in Δ_i models a possible assessment of potential certainties – more informally put, a possible way of *seeing* things, which leads to our being more or less confident of potential certainties in a particular manner. In turn, the perspective encoded in each P_j can be seen to spring from a number of particular circumstances in which the agent can be situated – circumstances that add to the constitution of the more encompassing context correlated with Δ_i .

A number of features of the present strategy are worth stressing. First, the model trades on well-known attempts to refine standard Bayesianism.²² Within such refinements,

¹⁹ Still, we might also demand that agents who find themselves holding inconsistent sets of acceptances at t should seek to eliminate the conflicting assumptions through one or more contextual contractions; I shall address the topic of contractions very briefly in the next section.

²⁰ Thanks to Paul Pedersen for pointing out a problem within a previous formulation of this definition.

²¹ Among other things, by proceeding thus we obtain that an expansion of A_i by elements that the agent judges to be epistemically possible in context i is a monotonic operation. Cf. Levi (2004), ch. 1, for a defense of this approach on conditionals and modal claims – within a rather different framework.

an agent is typically credited with the possibility of assigning *intervals* of probability measures to her uncertainties – where intervals, in turn, can be rendered as convex sets of probability functions. The present model can be seen, at least in part, as an attempt to extend this basic idea so as to take care of the need to distinguish between acceptances and full beliefs. In addition, we obtain a straightforward distinction between different senses of probability 1, insofar as contextual acceptances can be said to have probability 1 without thereby being full beliefs.²³ In this way, we stress the idea that there is a peculiar continuity between an agent’s state of certainty and doubt. Finally, in the next section I hope to show that, by letting contexts be associated with sets of probability functions in the manner just suggested, we obtain a neat way of tracking how contexts themselves change.

Before proceeding any further, let me consider two potential objections. First, let me address a concern related to my use of a probabilistic framework.²⁴ Someone might contend that we cannot have it both ways: either the model conceives of acceptance sets as “standards for serious possibilities” (SSP), as I suggested on p. 6, or the model takes probabilities to be basic, but the two things cannot be true at the same time. The reason (so the objection goes) is that SSP, to be such, need to be conceptually *prior* to any probability

²² Cf. the previous footnote. For some recent discussions on the advantages and problems of crediting agents with sets of probability measures see Elga (ms), White (2008), or Sturgeon (2008a, 2008b), to mention a few. I shall not enter into this debate here.

²³ It might be fruitful to compare the present attempt to distinguish different types of probability 1 with alternative proposals we find in the literature. Consider, for example, van Fraassen’s procedure in (1995). In that paper van Fraassen suggests a primitive notion of conditional probability, which is meant to take care of cases in which the condition has measure 0; in van Fraassen’s work, probabilities are applied to propositions as sets of points. His primitive notion of conditional probability then helps him define a system of nested belief cores K_i ; all the K_i s are sets with probability 1, and the inner belief core (if there is some core at all) is the one which intuitively carries stronger information. According to his terminology, the larger set with probability 1 is the set of *full beliefs* properly speaking. As we can see, one of van Fraassen’s main concerns is to make room for small increments of information, any of which will nonetheless have probability 0. In this sense, the goals of his proposal in (1995) and the goals of the present paper overlap only partially. Alternative systems that share at least some of van Fraassen’s motivations can be found in Arló Costa (2001), or Arló Costa and Parikh (2005), to mention a few. Notice that we can certainly combine these approaches with the model I am developing here, by letting the epistemic state of an agent be represented by a convex set of 2-*place* personal probability measures (defined in the spirit of systems with primitive conditional probabilities). Incidentally, this would be a way of recovering, for a probabilistic framework, van Fraassen’s own intuitions concerning the need to distinguish between believing and accepting, which is absent from his 1995 paper (although, of course, the way I am construing the belief/acceptance distinction is not faithful to van Fraassen’s own terminology in (1989)).

²⁴ Thank you to an anonymous referee for the *FEW 2008* for pressing this issue.

measure, as SSP are supposed to tell us which probability functions are admissible in the first place; admissible probability functions are then typically defined *conditional* on the agent's SSP.

The answer to this concern has already been anticipated at the beginning of the present section. In a nutshell, I take it that agents can intuitively be credited with having (among many other things):

- (a) A bunch of certainties that are not relative to particular circumstances or to the prior adoption of a particular point of view.
- (b) An array of competing ways of assessing how likely things appear to be, depending on the circumstances. Or, in other words, an array of competing ways of being more or less uncertain about the world.
- (c) A bunch of “conditional” certainties, so to speak – or certainties that are relative to the circumstances under which the agent's reasoning or acting may be taking place.

The array of rival assessments of potential certainties can be taken to presuppose a prior set of (conditional and unconditional) assumptions, whereas, conversely, the bunch of conditional and unconditional assumptions of the agent can be understood as the result of the agent's having the particular array of probability assessments she has. Typical agents, however, need not be aware of their possessing all these elements simultaneously, or of the essential inter-definability of (a), (b) and (c). As I have already suggested, quite often certainties *seem* more basic to the agent, which is precisely what licenses the theoretician to describe them as SSP in the first place (as I have done in my informal description of level (A), at the intuitive level, in section 3). This does not mean that such certainties should be conceived of as *ontologically* prior.

In short, let me emphasize that I do not assume either certainties or probability assessments to have ontological priority over each other. Notwithstanding, I have taken probabilities to be more basic *at the representation level*, for pragmatic reasons; as I can see it, in this way we obtain a simpler (and, to some extent, more elegant) epistemic model, which helps us arrive at K , M and the multiple A_i s with the aid of a single set of functions.²⁵

²⁵ Incidentally, notice that within the present framework probability 0 and 1 collapses with (epistemic) impossibility and necessity, respectively (in each context). In other words, unlike a model genuinely based on SSP, here we cannot allow that a sentence of L be seriously possible without receiving positive probability by

Consider now a second concern. Suppose that the agent can conceive of circumstances in which she would acknowledge the possibility of *not-p*. Doesn't this mean that *not-p* is possible for her, *simpliciter*? To put it differently, it could be objected that if the agent can represent to herself the possibility that such circumstances occur, then she can also represent to herself the possibility that *not-p* be the case, and hence there is no context in which *p* could be rendered as a contextual certainty. Conversely, if the agent is unable to represent to herself the possibility that such circumstances occur, that means that *p* is a full belief, rather than a mere contextual acceptance.

Against this contention, notice that contexts are *not* assumed to be describable with the (sole) aid of *L*. Thus, the correctness of a conditional such as “if I were to find myself reasoning under circumstances *c*, then *not-p* would be possible for me” cannot be evaluated in any context within the present framework – anymore than we can evaluate other modalities or counterfactuals: they are all *derivative* statements that would require an additional meta-theoretic apparatus in order to be formulated in the first place, and whose legitimacy (or lack thereof) would get exactly determined by the first order model, i.e., by the structure of Δ , and hence by the resulting *M* and *K*. Among other things, recall that contexts are partly defined by their set of epistemic assumptions – but only *partly* so. To say that an agent is currently situated in a particular context *c* means not just that she is situated in a concrete spatio-temporal location, but also that she has particular goals in mind, particular fears, particular worries, etc. Many of such elements are not appropriately rendered as the objects of epistemic attitudes – at least not in a primitive sense, although they might be so, again, in a *derivative* sense. In other words, descriptions of the potential circumstances an agent finds relevant at a particular time are not possible arguments of the functions in Δ ; if needed, we should attempt to reconstruct them *out of set* Δ - at a meta-theoretic level.²⁶

some function in Δ . This might lead to unwanted results in some cases; if dealing with such cases becomes important, I suggest enriching the present model along the lines discussed in footnote 23.

²⁶ I am indebted to Bernhard Nickel for his comments on this problem. A more complete treatment of this objection would require a full-fledged account of conditional statements, in which I cannot enter here. In particular, we should distinguish carefully between: (a) standard belief contravening conditionals in which the agent argues for the sake of the argument; and (b) counterfactual conditionals such that the antecedent describes circumstances that differ from the ones in which the agent is currently operating. In case (a) we need to reflect on how A_i gets modified for the sake of the argument, whereas in case (b) we need to shift our

5. Voluntary epistemic changes

In this section I shall explore how the model behaves at the time of representing epistemic changes. Let me begin by suggesting a few terminological conventions. I shall say that a contextual expansion takes place if, as a result of reasoning within context i , the agent ends up holding a set of acceptances A_j such that $A_i \subset A_j$, whereas I shall say that the agent performs a contextual contraction if, as a result of reasoning within context i , she ends up endorsing $A_j \subset A_i$. More generally, I shall say that *context i* changes if, as a result of reasoning within i the agent arrives at a set of acceptances that differs from A_i in some respect. Furthermore, I shall say that the very *space of contexts* changes if, as a result of a shift in some context, the agent ends up having more (or less) theories included in M – in other words, she ends up having more (or less) relevant contexts than before, new contexts may have been created, and older contexts may have merged together. As we shall see, changes in particular contexts need not modify the *space* of contexts; in addition, a set Δ_i may change without thereby provoking a substantial shift in context i – in the sense that the agent may arrive at a different set Δ_j which is nevertheless still correlated with A_i . In what follows I shall focus on describing *how expansions and contractions affect contexts*. The representation strategy will be compatible with several different rationales for coming to accept a particular statement α , as well as with many different specific suggestions on how to choose which additional statements to remove from a set in case we want to contract by α .

In (Cresto, 2006, 2008a) I have argued that conscious, voluntarily implemented expansions can be paradigmatically illustrated with instances of *inferences to the best explanation* (IBE), and I have also argued that the concept of IBE is best elucidated with the aid of some brand of cognitive decision theory.²⁷ In other words, I have suggested that we conceive of IBE as a decision theoretic exercise, in which we focus on the epistemic

attention to other sets A_j (*without* engaging in revision, not even for the sake of the argument). A careful exploration of the relation between (a) and (b) will be left for future work.

²⁷ The best-known approach to cognitive decision theory is found in Levi (1980); *cf.* also Maher (1993), or van Fraassen (1989, 2002).

gain we are able to obtain. According to this framework, agents come to accept best explanations *because they think that best explanations are worth the risk*. Indeed, agents risk being wrong – they risk accepting a false hypothesis – but taking the risk may be rational if the gain in overall understanding is high enough. IBE so conceived requires that we make reference to several kinds of contextual indices (such as caution thresholds, or contextual weights for the several dimensions that compose an epistemic utility function).^{28, 29} In addition to such indices, if our discussion from section 3 is on the right track, different instances of research may advance on the basis of different sets of epistemic assumptions. A similar analysis can be assumed to hold for voluntary expansions in general, regardless of whether we aim at the acceptance of a best explanatory hypothesis. From this perspective, *voluntarily coming to accept (hypotheses or statements) is a context-dependent activity*.

In a nutshell, the story may go like this. Consider the possibility of expanding acceptance set T_b with sentence α consistent with T_b , where T_b is, as usual, determined by a particular set Δ_b correlated with context b . (Cf. the Appendix for an illustration of this process). Let me assume that the agent can define the relevant parameters of a cognitive decision theoretic problem, including the identification of a suitable set of options for acceptance (which contains α), a suitable epistemic utility function, and perhaps a particular acceptance threshold (depending on the details of the particular brand of cognitive decision theory we favor). Before building the set of possible options – for instance, a set of explanatory hypotheses, if we are dealing with an IBE – the agent might feel compelled to seek for new evidence, which might lead her to update all measures in Δ_b through Bayesian conditionalization.³⁰ Next, the agent can use the chosen epistemic utility function to calculate the expected epistemic utility of the relevant sentences of L (i.e., of the

²⁸ In Cresto (2008b) I have argued that an adequate concept of epistemic utility should take into account the *virtuosity* of the basic explanatory hypotheses, which requires our paying attention to features such as simplicity, unification power, fertility, accuracy, or predictive force, among others.

²⁹ A word of caution. In certain discussions the expression “epistemic value” is used to refer to how much agents praise *particular judgments of personal probabilities*; epistemic values so conceived are usually elicited by means of so-called *scoring rules*. It should be clear that this is not the concept of epistemic utility that is at stake here.

³⁰ Cf. footnote 35, in section 6, for a further refinement of this claim.

options as determined by the decision theoretic problem), *for each probability function in Δ_b* . Then, for every P_i in Δ_b , if the expected epistemic utility of α , as calculated on the basis of P_i , is maximum and above the contextual threshold (in case there is a threshold at all), the theory recommends that $P_i(\alpha)$ be updated to 1.³¹ Notice that, depending on the case, the probability of α may end up being updated by all, none, or part of the elements of Δ_b . In the last case we obtain not only a contextual shift, but a change in the map of relevant contexts. Thus, by focusing on the way set Δ is modified, we have a way of capturing how contexts themselves evolve.³²

Let me consider now in a more detailed fashion the formal structure of this suggestion. In what follows I shall not be concerned with the development of an epistemic utility function (a task I have already pursued in previous work), and, in general, I shall not address the problem of how to build an adequate cognitive decision theory, although I shall assume that some such theory is possible. More generally, I shall assume that, regardless of the details, there is *some* account we can use at the time of deciding whether a given

³¹ It might be contended that allowing probability functions to be modified in the way I am advocating here commit us to a violation of so-called Bayesian conditionalization, and that we are therefore in trouble. In this paper I cannot enter into a detailed discussion of this topic, but, nonetheless, let me state very briefly a few considerations. To begin with, it is not clear whether the probability changes required by the present model constitute *violations* of the Bayesian conditionalization principle (what we might call “anti-Bayesian” shifts), or whether they are merely “non-Bayesian,” in the sense that the Bayesian conditionalization principle does not apply – insofar as the model does not recommend changes that takes place *as a result of* acquiring new evidence. (This, however, does not mean to say that empirical evidence does not play any role in acceptance, of course, but only that the decision theoretic exercise I am considering here takes place well after gathering the evidence and well after using such evidence to update probabilities in a Bayesian way). Against this line of reasoning, it might be suggested that *any* anti-Bayesian shift (according to the present terminology) can be decomposed into a Bayesian and a non-Bayesian step, and hence, to the extent that anti-Bayesian changes are irrational, so are non-Bayesian ones. In any case, I tend to think that the claim that all anti-Bayesian shifts are irrational is controversial and should not be accepted without substantial argumentation. In particular, let me recall here that arguments that appeal to diachronic Dutch Books have been contested on several occasions; cf. for example Levi (1987, 2002); Maher (1992); or Howson and Urbach (1993), pp. 99 and ff., among others.

³² It might be objected that it is just as rational to use a different rule, according to which we update the members of Δ_i with a given hypothesis H iff its expected epistemic utility is highest for *all* such members (thank you to Bernhard Nickel for raising this objection). However, the proposal I currently favor is the right way to go *given what each P_i is meant to represent*. Recall that each P_i stands for a particular way of conceiving of how *uncertain* things are, and that each P_i encodes the perspective the agent adopts under a particular cluster of circumstances. Moreover, the circumstances that shape each P_i add to the (possibly larger) context to which P_i belongs, which is then constituted, among other things, by a cluster of the different possible circumstances that correspond to different probability functions. If this is so, it seems just natural to say that, if adopting the perspective encoded by P_i leads to our giving maximum expected epistemic utility to H , then H should be accepted under the circumstances that correspond to P_i – regardless of the behavior of other probability distributions in the context.

probability function licenses the acceptance of a particular statement. Rather than focusing on what makes a statement *acceptable*, let me focus on how to proceed once we agree that our coming to accept the statement is legitimate, on the basis of a particular prior probability assignment – in short, let me concentrate on how to update acceptance sets and contexts (and, eventually – as we shall see in the next section – full beliefs).

Let $\Delta_{X,t}$ represent agent X 's epistemic state at t , as usual. I shall say that the shift from $\Delta_{X,t}$ to $\Delta'_{X,t'}$ is a contextual expansion of T_b by which agent X comes to accept α at t' in a particular context a (for α consistent with T_b) if and only if:

- 1) $\Delta'_{X,t'}$ is a convex set of probability measures.
- 2) There is some set $\Delta_a \subseteq \Delta'_{X,t'}$ (correlated with $T_a \subseteq M'_{X,t'}$)³³ such that for every P_n in Δ_a :
 - (a) $P_n(\alpha)=1$; and
 - (b) for any sentence γ , $P_n(\gamma) = P_m(\gamma/\alpha)$, for some P_m in $\Delta_b \subseteq \Delta_{X,t}$ correlated with $T_b \subseteq M_{X,t}$.
- 3) If there exists some $\Delta_s \subseteq \Delta_{X,t}$ correlated with a theory in $M_{X,t}$ that entails α , then $\Delta_b \neq \Delta_s$.
- 4) For all P_i in $\Delta_{X,t}$: either
 - (a) there is some P_n in Δ_a such that, for any sentence γ , $P_n(\gamma) = P_i(\gamma/\alpha)$; or
 - (b) $P_i \in \Delta'_{X,t'}$.
- 5) For all P_j in $\Delta'_{X,t'}$: either
 - (a) P_j was in $\Delta_{X,t}$; or
 - (b) $P_j \in \Delta_a$; or
 - (c) P_j is a mixture of elements satisfying (4.a) or (4.b).

(For an illustration of how this definition works in a concrete example, see the Appendix)

³³ The sign of improper inclusion is meant to acknowledge the possibility that $\Delta_{X,t}$ be actually a singleton. Incidentally, recall that, according to the notation I am using here, for any sub-index i , Δ_i is not to be thought of as an arbitrary subset of probability functions of Δ , but as the subset of Δ that can be correlated with context i , in the manner explained in the previous section.

Let us examine some consequences of this definition. First of all, it is straightforward from (2) that $T_a = Cn(T_b \cup \{\alpha\})$, as desired (where ‘ Cn ’ is, as usual, the Tarskian operator of logical consequence). In other words, now there is at least some function P_j in $\Delta'_{X,t}$ such that both $P_j(\alpha)=1$ and $P_j(T_b)=1$, and hence $f(P_j)=Cn(T_b \cup \{\alpha\})$.

Second, by (3) the definition entails that at least some probability function in Δ_b has been updated. This guarantees that we do not use the expression “ X comes to accept α at t ” in a particular context” just because α was already in some acceptance set before t – had there been a superset of T_b containing α already in $M_{X,t}$. This result also guarantees that, in case there existed Δ_d correlated with $Cn(T_b \cup \{\alpha\})$ in $\Delta_{X,t}$, then $\Delta_d \subseteq \Delta_a$, in which case a fusion of contexts would be obtained. In other words, in order to be fulfilled this definition requires that some real probability change occur, and that some acceptance set be actually enlarged, regardless of whether α had already been in M ; on the other hand, if α is a full belief of the agent at t , the definition cannot be satisfied: α needs to be, if not a full doubt, at least a contextual doubt.

In any case, our definition does not require that *all* functions in Δ_b be updated. Thus, Δ_b may end up split into two: part of its members may continue to be correlated with T_b (together with additional probability functions in order to satisfy convexity), while others will be associated with $Cn(T_b \cup \{\alpha\})$. Let $\mathbf{T}_M = \mathbf{T} \cap \wp^M$ be the set of theories (i.e., the set of acceptance sets) included in M . Then, depending on the case, we may obtain that the new set of theories \mathbf{T}'_M is either $\mathbf{T}_M \cup \{Cn(T_b \cup \{\alpha\})\}$ or the more economical $(\mathbf{T}_M \cup \{Cn(T_b \cup \{\alpha\})\}) \setminus \{T_b\}$. This is as it should be, if we want to allow for the possibility that the very space of contexts change. As I have already pointed out, the modification of the space of contexts should be distinguished from a mere shift in the set of probabilities *correlated with* some context, as well as from the evolution of particular contexts.

Finally, clause (4) guarantees that there are no unnecessary losses of former probability functions in the new set $\Delta'_{X,t}$, whereas by clause (5) the change from $\Delta_{X,t}$ to $\Delta'_{X,t}$ constitutes the smaller shift that fulfills (1)-(4).

It should be noticed that the aforementioned definition does not allow α to become an element of acceptance sets *properly included in* T_b . This is meant to reflect the fact that

the agent is not actually reasoning from within more cautious contexts – *i.e.*, from contexts in which she only takes for granted proper subsets of T_b .

To sum up, a voluntary expansion can be represented as the conscious implementation of a particular change in set Δ ; more precisely, a contextual expansion by α takes place when some probability function correlated with a particular context is voluntarily shifted so as to assign probability 1 to α , in which case α becomes a new contextual acceptance.

Contextual contractions proceed along similar lines, although, as is natural, in this case there will be several additional problems to take care of. A detailed analysis of contextual contractions will be left for future work. In particular, here I shall not examine the mechanisms that may lead to the decision to remove some sentence α from a specific theory.³⁴ Let me just suggest that a change from $\Delta_{X,t}$ to $\Delta'_{X,t'}$ is a *contextual contraction* by α at t' for a particular context (by which we mean to say that X ceases to accept α at t' , in that particular context) if and only if:

- 1) $\Delta'_{X,t'}$ is a convex set of probability measures.
- 2) There is some set $\Delta_a \subseteq \Delta_{X,t}$ correlated with $T_a \subseteq M_{X,t}$, such that $\alpha \in T_a$ and $\Delta_a \not\subseteq \Delta'_{X,t'}$
- 3) Suppose T_b is the largest theory properly included in T_a , and suppose $\alpha \notin T_b$. Then $\Delta'_{X,t'} = \Delta_{X,t} \setminus \Delta_a$.
- 4) Suppose there is no largest theory T_b properly included in T_a such that $\alpha \notin T_b$. Then there is $\Delta'_a \subseteq \Delta'_{X,t'}$ such that:
 - a) For every P_k in Δ'_a , $P_k(\alpha)$ is in some interval $(x,y) \subseteq (0,1)$.
 - b) For every P_i in Δ'_a and every sentence γ probabilistically independent of α (as far as X is concerned) there is some P_j in Δ_a such that $P_i(\gamma) = P_j(\gamma)$.
 - c) For every P_j in Δ_a and every sentence γ probabilistically independent of α (as far as X is concerned) there is some P_i in Δ'_a such that $P_i(\gamma) = P_j(\gamma)$.

Of course, this does not help us to uniquely determine the new set $\Delta'_{X,t'}$. In particular, in (4) I have not indicated how to proceed with sentences of L that the agent does *not* deem to be

³⁴ As mentioned briefly in section 4, we can identify a possible motivation to contract when, by reasoning within the decision theoretic framework described above, we end up having incompatible sets of acceptances.

probabilistically independent of α . Different approaches could be adopted here, which would lead to different instructions as to how to fix Δ'_a (and hence $\Delta'_{X,t}$). As before, if Δ'_a ends up being correlated with a theory that was already in $M_{X,t}$, we will have a fusion of contexts.

6. Involuntary epistemic changes

I have suggested that voluntary epistemic expansions are essentially *contextual* expansions. The way I see it, at the time of engaging in research and reflecting on the particular “gain” we obtain from a given hypothesis or statement (for instance, a particular explanatory relief), an agent does not consider the advantages or disadvantages of *coming to believe*, for all contexts, that the statement is true. This is not a goal, precisely because it cannot be consciously implemented: coming to believe that something is the case is not voluntary.

This is not to say, however, that agents are never able to modify their stock of full beliefs. In the first place, *in addition* to the contextual expansions that I have described so far, the present model can well allow for spontaneous, involuntary expansions of K , which will affect every acceptance set $A_i \subseteq M$. The acquisition of new perceptual beliefs constitutes a paradigmatic example of this situation.³⁵

Second, and even more interesting, notice that, under certain conditions, successive contextual credal shifts might lead to changes in the agent’s set of full beliefs. More precisely, if successive shifts in Δ are such that statement α is finally given probability 1 by all functions, we will obtain a bona fide belief expansion by α – that is, a change from K to $K' = Cn(K \cup \{\alpha\})$. Thus, within this framework many elements in $M \setminus K$ will eventually “leak” into K , but this is actually a side effect of performing multiple adjustments in the total lot of the agent’s acceptance sets. In this sense, expansions of K are not under the agent’s control, even though they can be the by-products of voluntary contextual

³⁵ Related to this, it might be argued that the gathering of *new evidence* always results in involuntary expansions of K of this sort – in which case *all* elements of $\Delta_{X,t}$ (and not just those correlated with the context in which the particular inquiry is taking place) would be updated accordingly, in a Bayesian way.

expansions: the reason is that each contextual change does not have *the goal* of modifying her set of beliefs.³⁶

Formally speaking, we can say that agent X comes to believe α at t_n (for α compatible with $M_{X,t}$) if and only if:

Either

(A) There is a spontaneous change from $\Delta^m_{X,tm}$ to $\Delta^n_{X,tn}$ such that:

- 1) For some P_i in $\Delta^m_{X,tm}$, $P_i(\alpha) \neq 1$.
- 2) For every P_j in $\Delta^m_{X,tm}$ and all γ , there is some P_k in $\Delta^n_{X,tn}$ such that $P_k(\alpha)=1$, and $P_k(\gamma) = P_j(\gamma/\alpha)$.
- 3) Nothing else is in $\Delta^n_{X,tn}$;

Or:

(B) There is a sequence of consciously implemented contextual expansions $\Delta^1_{X,t1}, \dots, \Delta^m_{X,tm}, \Delta^n_{X,tn}$ (as defined in section 5), such that:

- 1) For some P_i in $\Delta^1_{X,t1}$, $P_i(\alpha) \neq 1$; and
- 2) For all P_j in $\Delta^n_{X,tn}$, $P_j(\alpha)=1$.

A consequence of this approach is that agents can be, as a matter of fact, justified in coming to believe α (say, if the step from $\Delta^m_{X,tm}$ to $\Delta^n_{X,tn}$ is justified in the manner explained in section 5, or perhaps if the involuntary expansion was the result of a reliable process), though they can never be justified in *seeking to believe* α .

A similar story could in principle be told for contractions, although I shall not enter into the details here. As with acceptances, full doubts are not voluntary (as many critics of skepticism have suggested before), but we can stumble upon them, so to speak, after a sequence of contextual contractions.³⁷ In addition, the model also allows for the occurrence of *spontaneous doubts*: namely, the model allows for the occurrence of shifts in Δ that the

³⁶ Notice that, according to the definitions suggested in previous sections, in order to obtain an expansion for all contexts we have to take into account what I have dubbed “the minimal context” as well – which is correlated with the very same K . Now, if the model allows that we reason on the basis of K , shouldn’t this mean that we can aim for *belief* changes? The answer is clearly ‘no’: if we accept that p when reasoning on the basis of K , we need not succeed in *changing* K ; typically, what we obtain is a new acceptance set that properly includes K .

³⁷ Once again, in order to remove a sentence from all contexts – so that it becomes a full doubt – we need to reason on the basis of K as well. Cf. the previous footnote.

agent has not consciously chosen to implement, and which have the effect of erasing a given sentence from every acceptance set in M .

7. Conclusions

Let me summarize briefly what we have achieved. In section 1 I suggested that we use the distinction between belief and acceptance to obtain an epistemological picture in which we could acknowledge the existence of voluntary and involuntary epistemic changes, as well as the presence of contextual and pragmatic factors. In section 2 I identified different levels of analysis, and stressed the fact that the paper would proceed at a *representation level*, by capturing some phenomena that were already intuitively clear at the level of the agent's *real* epistemic state, while being neutral with regard to the ontological problem on beliefs and acceptances. In section 3 I presented the intuitive distinction between contextual epistemic assumptions and epistemic assumptions held in all contexts; I coined the expression “acceptances” as a term of art to refer to the set of epistemic assumptions of a particular context, and “full beliefs” to refer to the set of epistemic assumptions held in all relevant contexts of action and deliberation (as far as the agent is concerned). I also argued that what counts as a relevant context for an agent at a particular time depends in part on which sets of circumstances the agent can conceive of and deem possible to affect her reasoning; in this sense, determining which contexts are relevant is up to each agent.

In sections 4, 5 and 6 I developed the bulk of my formal proposal. I suggested that an agent's epistemic state be modeled by a convex set Δ of probability functions P_i over the sentences of a regimented language L . Each P_i is meant to represent a particular assessment of potential certainties; such assessment can be seen to spring from a number of particular circumstances in which the agent can be situated – circumstances that add to the constitution of what the agents takes to be some relevant context. I also defined a function from Δ to \mathbf{T} (the set of theories of L), which mapped each probability measure in Δ to the largest theory whose elements are assigned probability 1 by that measure. A set of contextual acceptances is then a theory that is selected by some member of Δ , and contexts can be said to be correlated with particular subsets of Δ - that is, with those subsets that

contain all probability functions that pick out the same theory. Set K , the set of full beliefs, is defined as the theory that receives probability 1 by all functions in Δ , and hence is the intersection of all acceptance sets.

Next I focused on epistemic changes. A contextual expansion by α takes place when some probability function correlated with a particular context is shifted so as to assign probability 1 to α ; hence there is some T_i in the agent's older epistemic state, and some context j in her new set of contexts, such that the agent now counts with a new acceptance set $T_j = Cn(T_i \cup \{\alpha\})$. Under this perspective, voluntarily coming to accept a statement expands M (the set of statements accepted at *some* context), but not necessarily K . In addition, it might happen that, by means of reasoning within context i , the agent comes to realize that she is not ready to take α for granted under the whole array of circumstances that contributed to the identification of context i in the first place, but only under a more restricted set of circumstances – in which case *a new context* starts to be considered relevant by the agent. On the other hand, a contextual contraction by α occurs if some measures in Δ are shifted in such a way that they no longer give probability 1 to α .

In section 6 I also stated that the model allows for spontaneous changes in K . In addition, I suggested that, under certain conditions, successive contextual credal shifts may lead to changes in K . In particular, if successive shifts in Δ are such that statement α is finally given probability 1 by all functions, we obtain a bona fide belief expansion by α . Thus, within this framework many elements in $M \setminus K$ eventually enter into K , as a side effect of performing multiple adjustments in \mathbf{T}_M (the set of all accepted theories). Analogously, we can arrive at full doubts after a sequence of contextual contractions.

In short, by distinguishing between beliefs and acceptances the model succeeds in acknowledging the existence of voluntary and involuntary aspects of our epistemic life, while keeping both aspects integrated within the same account. In addition, as promised, the model presents an epistemological picture in which pragmatic factors play a crucial role,³⁸ and it makes room for the idea that our voluntary epistemic life proceeds in a

³⁸ For different attempts to reflect on the presence of pragmatic factors in epistemology cf. Fantl and McGrath (2002), Hawthorne (2004), Stanley (2005), or Hawthorne and Stanley (forthcoming), among others.

thoroughly contextual way. In other words, the model delivers just what we hoped to obtain. As for the more technical features, by allowing for vague probability assignments the present framework attempts to overcome a well known difficulty of standard Bayesianism; by the same token, the strategy of letting epistemic states be represented by *sets* of functions helps us distinguish probability 1 from full belief, while by making it a *probabilistic* model, we showed how to treat probability and full belief in a unified way – thus emphasizing the idea that there is a peculiar continuity between doubts and certainties. Finally, the chosen setting enables us to represent how the overall map of contexts evolve as a result of particular credal shifts.

Appendix

Suppose Jill Jones is a biologist interested in immunology. Lately, her research has begun to focus on a long-lasting perplexity of the scientific community: why is it that, during pregnancy, the woman's immune system does not attack the growing fetus, despite the fact that half of its genes are alien to the mother? Until recently, the most popular hypothesis suggested that the placenta somehow acted as a mechanical barrier that prevented T-cells from harming the fetus – but, as a matter of fact, this claim raised as many worries as it helped to answer. Last year Jill joined a team whose research line dealt with exploring the consequences of some novel suggestions about the growth of tumors. Pretty soon a parallel between the two processes became apparent to her, as well as to many of her colleagues. They performed several experiments with mice, and obtained results consistent with the idea that progesterone stimulated the production of galectin-1 (Gal-1), an immunoregulatory glycan-binding protein (cf. Blois *et al.*, 2007). In addition, independent tests were consistent with the idea that Gal-1 caused the induction of tolerogenic dendritic cells; as opposed to “regular” dendritic cells, tolerogenic dendritic cells have lost their ability to activate T-lymphocytes, and they promote the expansion of interleukin-10 (IL-10)–secreting regulatory T cells (which block the immune attack); cf. (Toscano *et al.*, 2007). In the light of this, now Jill and her colleagues are aware of a novel possible explanation for fetomaternal tolerance – to wit, that during normal pregnancy, progesterone stimulates the production of Gal-1, which induces tolerogenic dendritic cells (a mechanism already postulated in tumor growth), ultimately suppressing T-cell activity against the fetus.

When doing research, Jill finds herself convinced that the explanation for fetomaternal tolerance lies either in the barrier hypothesis, or in the one that stresses the role of Gal-1; she feels *certain* about this disjunction. In more mundane contexts, however (say, when not pressed by the urge to obtain results relevant to her current research line) she is not so convinced – actually, in more mundane contexts she tends to be wary about fully assenting to the truth of general explanations, or even to the truth of (non-tautological) disjunctions of explanations.

On a different line, it turns out that Jill is throwing a party tomorrow. She goes to the supermarket to buy some groceries, and finds herself taking for granted that John will

come to her party. When she is back at the lab, someone asks her whether John is coming tomorrow; she reflects for a second and says, “oh, you know, you can never be sure of what people will do. If I had to guess, I’d say, *very* probably yes”. In other words, she is not certain. Intuitively, however, she has not changed her mind – she has only changed her context of action and deliberation. Actually, while she is at the lab she tends to be wary about fully assenting to statements that describe the occurrence of future events that involve human planning of some sort.

Consider now an extremely simplified model of Jill’s epistemic state at t_0 . Assume the set of sentences of L is built recursively out of atomic formulae $p, q, r_1, \dots, r_n, s, t$ and u , where the intuitive translations are as follows:

p : I am Jill Jones.

q : During normal pregnancy, the mother’s immune system does not attack the fetus.

s : John will come to my party tomorrow.

t : During normal pregnancy, the placenta acts as a mechanical barrier that blocks the mother’s immune attack on the fetus.

u : During normal pregnancy, progesterone stimulates the production of Gal-1, which induces tolerogenic dendritic cells (a mechanism already postulated in tumor growth), ultimately suppressing T-cell activity against the fetus.

In addition, $r_1 \dots r_n$ stand for various descriptions of observed experimental results.

I shall assume that t and u are materially inconsistent with each other, and that they materially entail q ; t and u can be conceived of as rival explanations of fetomaternal tolerance, as found in normal pregnancy. At t_0 Jill holds probability measures P_1, P_2 and P_3 , where:

$$P_1(p)=P_1(q)=P_1(r_1)=\dots=P_1(r_n)=P_1(s)=1; P_1(t)=0.2; P_1(u)=0.6$$

$$P_2(p)=P_2(q)=P_2(r_1)=\dots=P_2(r_n)=1; P_2(s)=0.99; P_2(t)=0.4; P_2(u)=0.6$$

$$P_3(p)=P_3(q)=P_3(r_1)=\dots=P_3(r_n)=1; P_3(s)=0.99; P_3(t)=0.2; P_3(u)=0.8$$

Jill’s epistemic state at t_0 is then represented by the convex set built out of P_1, P_2 and P_3 ; let us call it Δ_{J,t_0} . We obtain three relevant contexts, correlated with the following subsets of Δ_{J,t_0} :

$$\Delta_a = \{P_1\}$$

$$\Delta_b = \{P_i; P_i=P_2, \text{ or } P_i=P_3, \text{ or } P_i \text{ is a mixture between } P_2 \text{ and } P_3\}$$

$$\Delta_k = \{P_i: P_i \text{ is a mixture between } P_l \text{ and } P_j, \text{ for all } P_j \in \Delta_b\}$$

(where k is the minimal context).

In other words, at t_0 Jill is *fully* certain of p, q, r_1, \dots, r_n , and their logical consequences: regardless of the circumstances, at t_0 she takes p, q, r_1, \dots, r_n and their consequences for granted. In addition, in context b she has accepted that the explanation for fetomaternal tolerance lies either in t or in u , whereas she is not certain about s ; in context a , by contrast, she accepts that s is the case, but she is not convinced of $t \vee u$.³⁹

Somewhat artificially, here I am assuming that in context a Jill has a single way of assessing how uncertain it is that the explanation for fetomaternal tolerance be either the barrier hypothesis, or the one linking progesterone to Gal-1 and to tolerogenic dendritic cells – which is rendered as $t \vee u$'s having a precise probability in a . Likewise, I am assuming, for the sake of simplicity, that in context b Jill has a single way of assessing how uncertain it is that John will come to the party tomorrow – which results, once again, in s 's having a precise probability in b . By contrast, in context b we find a variety of stances towards both t and u . Intuitively, each such assessment represents a more or less cautious attitude towards the two hypotheses – which might encode Jill's reflection on the different objective performances of several past tests.

Notice that at t_0 Jill is fully in doubt about what counts as a best explanation for q . Suppose she is at work right now; she discusses the problem of fetomaternal tolerance with her colleagues, and the question comes up: which explanation is better? In particular, is any of the two good enough so as to fully endorse it? Jill starts to reflect on the results of tests, on how probable she takes them to make each of the two hypotheses, on how much she would gain, epistemically speaking, by adopting one or the other as true (say, would she attain a better overall understanding of biological processes, for instance?) and she realizes that, all things considered, it seems just appropriate for her to adopt the hypothesis about the role of Gal-1 as true at the time of going on researching on cancer mechanisms – but not at the time of working on therapies for pregnancies at risk. This is tantamount to saying that Jill assesses whether she is in a position to arrive at a best explanation for q within her

³⁹ Recall that, within the present model, different possible circumstances are rendered as the same epistemic context if they lead the agent to embrace the same set of epistemic assumptions. Cf. p. 8 above.

current epistemic situation. She reasons within context b , and considers whether to expand by t , by u , or not expand at all, as the case may be.⁴⁰ We can reconstruct her inference to the best explanation along the following lines.

According to our example, the probability of t adopts all values in $[0.2, 0.4]$, and the probability of u adopts values in $[0.6, 0.8]$. Suppose Jill assesses the epistemic utility of both t and u in a 0-1 scale, and assigns epistemic utilities of 0.4 and 0.8 to them, respectively. (Notice that u clearly excels in both unification power and fertility, insofar as it helps to get a unified picture of two seemingly disparate phenomena – fetus and tumor growth – while it suggests a path of further tests that might ultimately lead to therapies to enhance both tumor survival and reduction of spontaneous abortions.)

Let me assume for the sake of concreteness that, according to our favorite theory, $EEU(h) = P(h)eu(h)$, where h is a legitimate option for expansion, $eu(h)$ stands for h 's epistemic utility, and $EEU(h)$ stands for h 's *expected* epistemic utility.⁴¹ Further, suppose that Jill is only ready to accept h in the context if $EEU(h)$ is maximal and above 0.5. Then we have,

For P_2 :

$$EEU(u)=(0.6)(0.8)=0.48 > EEU(t)=(0.4)(0.4)=0.16;$$

For P_3 :

$$EEU(u)=(0.8)(0.8)=0.64 > EEU(t)=(0.2)(0.4)=0.08$$

As 0.48 is below the chosen caution threshold, some probability functions in the context will be updated, but not all; hence, a new map of contexts will be obtained. More precisely, for all P_i in Δ_b such that $P_i(u) \in (0.625, 0.8]$, P_i will be replaced by P_3' , where

$$P_3'(p)=P_3'(q)=P_3'(r_1)=\dots=P_3'(r_n)=P_3'(u)=P_3'(-t)=1; P_3'(s)=0.99$$

(Of course, the situation would have been rather different if P_2 and P_3 assigned different values to statements probabilistically independent of t , such as s . But we need not enter into this complication here.)

In addition, Jill's new epistemic state will have functions P_2' and P_2'' such that:

⁴⁰ Incidentally, in the course of her research many evidential statements (such as the observational results of several tests) have already become full certainties of Jill; we can suppose that they have already helped update the probability of both t and u in every context, in a standard Bayesian manner.

⁴¹ A rationale for this suggestion can be found in the thought that, if we assume a hypothesis to be false, then no epistemic satisfaction can be obtained from it, and hence $eu(-h)$ is 0.

$$P_2'(p)=P_2'(q)=P_2'(r_1)=\dots=P_2'(r_n)=1; P_2'(s)=0.99; P_2'(t)=0.4; P_2'(u)=0.6;$$

and

$$P_2''(p)=P_2''(q)=P_2''(r_1)=\dots=P_2''(r_n)=1; P_2''(s)=0.99; P_2''(t)=0.375; P_2''(u)=0.625$$

Thus, at t_I Jill holds $\Delta'_{J,tI}$ with contexts a' , b' , b'' and k' , as follows:

$$\Delta_a' = \Delta_a = \{P_I\}$$

$$\Delta_b' = \{P_3'\}$$

$$\Delta_b'' = \{P_i: P_i=P_2', \text{ or } P_i=P_2'', \text{ or } P_i \text{ is a mixture between any of } P_2', P_2'' \text{ or } P_3'\}$$

$$\Delta_k' = \text{All remaining mixtures}$$

The idea is that Jill has just accepted u as a best explanation in context b' , whereas no expansion took place at context b'' . Intuitively, insofar as a split of contexts has just occurred, b' no longer exhausts all relevant possible circumstances in which Jill might find herself while doing research (as far as she is concerned), as we wanted. The structure of $\Delta'_{J,tI}$ tells us that the many circumstances in which she is ready to accept $t \vee u$ as true are not all alike with respect to u . Once b' and b'' are distinguished from each other, Jill can go on to define particular contextual parameters, such as particular caution thresholds, to keep on reasoning within each of the two contexts separately.

References

Arló Costa, H. (2001). "Bayesian Epistemology and Epistemic Conditionals: On the Status of the Export-Import Laws." *Journal of Philosophy*, 98: 555-598.

Arló Costa, H., and Parikh, R. (2005). "Conditional Probability and Defeasible Inference." *Journal of Philosophical Logic*, 34: 97-119.

Blois, S. M.; Ilarregui, J. M.; Tometten, M.; García, M.; Orsal, A. S.; Cordo-Russo, R.; Toscano, M. A.; Bianco, Germán A.; Kobelt, P.; Handjiski, B.; Tirado, I.; Market, U. R.; Klapp, B. F.; Poirier, F.; Szekeres-Bartho, J. L.; Rabinovich, G. A. & Arck, P. C. (2007). "A pivotal role for galectin-1 in fetomaternal tolerance." *Nature Medicine*, 13: 1450 – 1457.

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

Bilgrami, A. (2000). Is Truth a Goal of Inquiry? Rorty and Davidson on Truth. In R. Brandon (Ed.), *Rorty and his critics* (pp. 242-261). Oxford: Blackwell.

Bilgrami, A. (2004). Skepticism and Pragmatism. In D. McManus (Ed.), *Wittgenstein and scepticism* (pp. 56-75). London: Routledge.

Bratman, M. (1992). "Practical reasoning and acceptance in a context". *Mind*, 101: 1-15

Cohen, J. (1992). *An Essay on Belief and Acceptance*. Oxford: Clarendon Press.

Cohen, S. (2000). "Contextualism and Skepticism." *Philosophical Topics*, 10: 94-107.

Cresto, E. (2006). *Inferring to the Best Explanation: A Decision-Theoretic Approach*. Ph.D. Thesis, Columbia University, New York.

Cresto, E. (2008a). "In search of the best explanation about the nature of the gene: Avery on DNA and pneumococcal transformation". *Studies in History and Philosophy of Biological and Biomedical Sciences*, 39: 65-79.

Cresto, E. (2008b). "A model for structural changes of belief." *Studia Logica*, 88: 431–451.

Da Costa, N. and French, S. (2003). *Science and Partial Truth: A Unitary Approach to Models and Scientific Reasoning*. New York: Oxford University Press.

DeRose, K. (2002). "Assertion, Knowledge, and Context". *Philosophical Review*, 111: 167-203.

- Elga, A. (unpublished manuscript) 'Subjective Probabilities Should be Sharp'.
- Engel, P. (ed.) (2000). *Believing and Accepting*. Dordrecht: Kluwer.
- Fantl, J. and McGrath, M. (2002). "Evidence, pragmatics and justification." *The Philosophical Review*, 111: 67-94.
- Gärdenfors, P. (1988). *Knowledge in Flux – Modeling the Dynamics of Epistemic States*. Cambridge, Mass.: The MIT Press.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. and Stanley, J. (forthcoming). "Knowledge and Action". *The Journal of Philosophy*.
- Howson, C. and Urbach, P. (1993). *Scientific Reasoning: The Bayesian Approach*. Second edition. Chicago and La Salle, Illinois, Open Court. (First edition: 1989).
- Kaplan, M. (1996). *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.
- Lehrer, K. (2000). *Theory of Knowledge*. Boulder: Westview Press.
- Levi, I. (1974). "On Indeterminate Probabilities." *The Journal of Philosophy*, 71: 391-418.
- Levi, I. (1980). *The Enterprise of Knowledge*. Cambridge, Mass.: The MIT Press.
- Levi, I. (1987). "The Demons of Decision." *The Monist*, 70: 193-211.
- Levi, I. (1997). *The Covenant of Reason*. Cambridge: Cambridge University Press.
- Levi, I. (2002). "Money Pumps and Diachronic Dutch Books." *Philosophy of Science*, 69: S235-S247.
- Levi, I. (2004). *Mild Contraction: Evaluating Loss of Information due to Loss of Belief*. New York: Oxford University Press.
- Maher, P. (1992). "Diachronic Rationality." *Philosophy of Science*, 59: 120-141.
- Maher, P. (1993). *Betting on Theories*. Cambridge: Cambridge University Press.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Nozick, R. (1993). *The Nature of Rationality*. Princeton: Princeton University Press.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, Mass.: The MIT Press.

- Sturgeon, S. (2008a). "Reason and the Grain of Belief". *Noûs*, 42: 139-165.
- Sturgeon, S. (2008b). "Confidence and Coarse-Grained Attitudes." Manuscript. Available at <http://socrates.berkeley.edu/~fitelson/few/sturgeon.pdf>
- Toscano, M. A.; Bianco, G. A.; Ilarregui, J. M.; Croci, D. O.; Correale, J.; Hernandez, J. D.; Zwirner, N. W.; Poirier, F.; Riley, E. M.; Baum, L. G. & Rabinovich, G. A. (2007). "Differential glycosylation of T_H1, T_H2 and T_H-17 effector cells selectively regulates susceptibility to cell death." *Nature Immunology*, 8: 825 – 834.
- Tuomela, R. (2000). "Belief versus Acceptance." *Philosophical Explorations*, 2: 122-157.
- Van Fraassen, B. (1989). *Laws and Symmetry*. Oxford: Clarendon Press.
- Van Fraassen, B. (1995). "Fine grained option, probability, and the logic of full belief." *Journal of Philosophical Logic*, 24: 349-377.
- Van Fraassen, B. (2002). *The Empirical Stance*. New Haven: Yale University Press.
- White, R. (2008). "Evidential Symmetry and Mushy Credence". Manuscript. Available at <http://socrates.berkeley.edu/~fitelson/few/white.pdf>