

BELIEF AND CONTEXTUAL ACCEPTANCE

Eleonora Cresto

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

eleonora.cresto@gmail.com

[Published in *Synthese*, 177 (1), 2010, pp. 41-66]

ABSTRACT

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.

Keywords: Belief revision - Acceptance - Bayesianism - Cognitive decision theory - Formal epistemology

1. Introduction

In this paper I propose a strategy for modeling the epistemic state and epistemic changes of a particular agent at a given time, in the tradition of formal belief revision theories. The model seeks to

- (a) illuminate the extent to which there is room for pragmatic factors in epistemology;
- (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 shall argue that, insofar as these goals are fulfilled, we obtain a representation tool that preserves many pre-theoretical intuitions about what counts as suitable rationality constraints on our actual practices – and, in this sense, the model can be said to encode the basic guidelines of an adequate, normative epistemological theory. In addition, I hope to show that the account I offer here exhibits several technical advantages over alternative ways of proceeding. To carry out this project I shall rely partially on 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 paper is organized as follows. In section 2 I offer an example to motivate the analysis; in section 3 I draw some morals from the example, whereas in section 4 I suggest

a number of conceptual distinctions, on the basis of which I then build a formal framework. I develop the formal proposal in sections 5 to 9; section 5 deals with the basic structure of the model, sections 6 to 8 address voluntary expansions and contractions, and section 9 examines involuntary changes. Finally, in section 10 I present some conclusions.

2. Jill's case

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 – even though, 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-secreting regulatory T cells, which block the immune attack (cf. Toscano *et al.*, 2007). In the light of this, Jill and her colleagues are now 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.

Until now, Jill has been fully in doubt about what counts as a best explanation for the phenomenon of fetomaternal tolerance. 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 they 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.

On a different line, it turns out that Jill is throwing a party tomorrow. She goes to the supermarket to buy some groceries for her guests, 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 actually 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. As a matter of fact, 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.

I take it that the sort of doubts and certainties that Jill exhibits here, as well as her attempts to settle her mind, are representative of many of our usual epistemological practices, and hence it could be instructive to draw a few morals from her case. I shall undertake this task in the next section.

3. Some morals

(I) Agents sometimes seem to make epistemic decisions

There’s a well-entrenched philosophical tradition that insists that we cannot believe at will; beliefs are taken to be *involuntary*. However, sometimes agents make non-deductive inferences and seek to change their minds on the basis of the results of such inferences; they also seek to convince others through rational conversation. Take, for example, Jill’s story: after careful reflection, it seems perfectly rational for her to assent (voluntarily) to the truth of the hypothesis about the role of Gal-1 for fetomaternal tolerance, at least for the restricted set of circumstances under which she is considering assenting to its truth.

How can we reconcile these opposite intuitions? We might be tempted to think that part of the problem here is that we can adopt very different conceptions of belief, *ontologically speaking*. For example, if beliefs are mainly characterized as epistemic commitments, in the sense of (Levi 1980, or 1997), Humean-style involuntarism does not look too promising; the very idea of commitment embodies an irreducible normative element, and 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: dispositions can well be acquired (and maintained) without our willing this to happen.

Still, a clarification of our ontological assumptions on epistemic matters will not lead us too far. The reason is that we want *both* voluntarism and involuntarism to hold, albeit for different scenarios and different types of examples. We do not want a model that allows for the possibility that a given perceptual statement be treated as true [false] out of sheer will power; symmetrically, we do not want to be allowed to disregard the truth of a statement just because our gut feelings tell us to do so, if on careful reflection we have found very good reasons to adopt it as true.

In the light of this, we might want to elaborate a distinction between *believing* and *accepting* (hypotheses or propositions), in order to restore consistency and alleviate the tension. As is well known, different variants of this dichotomy have been proposed during recent decades, with the aim of solving very different problems; not surprisingly, a quick

look at the literature shows that there is no uniform way of understanding these concepts.¹ In section 4 I shall develop my own version of a belief-acceptance distinction that will help us reconcile voluntaristic and involuntaristic intuitions while also enabling us to account for other important features of Jill's case.

(II) Pragmatic factors are relevant at the time of making epistemic decisions

Traditionally, practical concerns were not deemed relevant for epistemic matters. However, on a close reflection they cannot be so easily eliminated.² When Jill considers which hypothesis to fully endorse, for example, it does not seem *rational* for her to adopt the one about the role of Gal-1 at the time of devising therapies for pregnant women, whereas she does feel entitled to endorse it at the time of researching for cancer mechanisms. This gives support to the thought that pragmatic factors actually affect our epistemic life: at the very least, they play a role when we assess the legitimacy of making various epistemic decisions – including the decision to withhold judgment, as the case may be.

(III) Even more generally, the circumstances under which an agent acts or reasons affect what she is ready to take for granted at that particular moment.

Recall that, while working at the lab – and even *before* seeking to assess which explanation is better – Jill takes for granted the truth of a particular disjunction of hypotheses; when she is not at work, by contrast, she does not feel equally convinced. This example suggests that at least part of what we happen to take for granted at the time of acting or reasoning typically depends on the circumstances under which our acting or reasoning takes place. In other words, the particular circumstances at stake (such as the type of audience we are addressing, or our degree of emotional involvement with the topic, as well as the various specific practical concerns we may have), lead us to proceed under the assumption that a number of ideas or propositions (I shall leave the terminology here intentionally vague) *are true*: as far as we are concerned, they are not open to discussion at the moment – which is of course consequential, in turn, for the type of epistemic decisions we will end up making in each case. To put it differently, part of the assumptions that an agent holds at any given time may *vary* with the circumstances. Other assumptions, by contrast, seem not to be so

¹ 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. From a different perspective, 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 – though, as we shall see, their understanding of 'context' is very different from the one that will be favored in this paper. Cf. also the articles in Engel (2000). For yet other proposals see Nozick (1993), Kaplan (1996), Tuomela (2000), or Da Costa and French (2003).

² For some recent attempts to reflect on the presence of pragmatic factors in epistemology cf. Fantl and McGrath (2002), Hawthorne (2004), Stanley (2005), or Hawthorne and Stanley (2008), among others. Their work, however, proceeds from a very different standpoint than the one that will be adopted in this paper.

restricted; for example, at present Jill does not conceive of any set of circumstances under which she may doubt her actually being Jill Jones.

A note of caution: typically, when arguing or deliberating agents also take 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. From now on I shall use the expression ‘*epistemic assumptions*’ to refer to those assumptions (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.³ For the most part, in what follows I shall not be concerned with assumptions that are not epistemic, in the sense just mentioned.

(IV) A shift of epistemic assumptions need not produce a genuine belief change.

Jill’s reluctance to take John’s coming to her party for granted when she is back at the lab points to the fact that agents typically hold a peculiar continuity between certainties and uncertainties: we feel *more or less uncertain* about a number of things under particular circumstances, and *not in doubt* about those very same things when we are in different scenarios. We can go back and forth from certainty to doubt, depending on the circumstances – which makes it very counter-intuitive to render the whole phenomenon as an actual epistemic change. Thus, a switch of epistemic assumptions due to a change of circumstances, in the sense of (III), does not seem equivalent to a bona fide change in view. For a different example, at the time of interacting with my neighbor I feel 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. It is not, however, that I have changed my mind: every time I speak to him, I do not take the possibility that he is not a college student seriously.

In the next section I shall rely on these considerations to develop a number of philosophical distinctions that will lie at the heart of the formal model I offer in sections 5 to 9.

4. Beliefs and contextual certainties

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. Thus, the examples discussed above allow us to say 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*,⁴ where contexts are partly defined by sets of epistemic assumptions. For pragmatic reasons, as this paper is focused on the structure of

³ In particular, credal probability judgments are not epistemic assumptions; I shall come back to this point below.

⁴ Unfortunately, the words ‘context’ and ‘contextual’ are heavily charged in the philosophical literature; see below for some caveats on the use of these terms.

epistemic states, I shall stipulate that if contexts i and j are constituted by the same sets of epistemic assumptions, then $i=j$. In addition, I shall say that an agent may hold ‘contextual doubts’ (doubts that are relative to one specific context but not to others), as much as ‘absolute doubts.’ I shall also assume that doubts can be probabilistically ordered; we will have much more to say about doubts in the next section.

Let me proceed now to offer a belief-acceptance distinction that helps us to systematize the ideas we have discussed so far. I shall speak of ‘acceptances’ to refer to the agent’s epistemic assumptions (or certainties) in some particular context, whereas the label ‘full beliefs’ (or ‘beliefs’, for short) will refer to epistemic assumptions held 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. Moreover, coming to accept is meant to be voluntary – but not so coming to believe, insofar as deliberating on epistemic matters is irremediably contextual. The way the present belief-acceptance distinction is supposed to articulate the voluntary-involuntary dichotomy will be apparent in further sections.

A number of precisions are in order:

- (1) When Jill is at the supermarket and about to buy food for her guests, she just takes for granted that John will come to her party; she is clearly not proceeding ‘for the sake of the argument’. In general, epistemic assumptions in the sense I am interested here lead agents *to act* in certain definite ways, which suggests that contexts in my sense are not to be confused with contexts of suppositional reasoning.⁵ Moreover, when we say that an epistemic decision has been made 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 adopts 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.
- (2) A context is not, in general, given by the sentences that the subject would be willing to assert; asserting that p might, but *need not*, be relevant to discovering the agent’s real epistemic assumptions.⁶ Even more generally, not every situation in which an agent appears to act *as if p* will be indicative of her being certain of p ; conversely, being certain of p will not always result in her acting *as if p*. Suppose I am trying to convince you of the truth of hypothesis h ; in order not to beg any questions, I pretend to be in suspense about h 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 one in which I am certain that h is the case: that’s why I bother offering an argument for h in the first place. (And notice, incidentally, that it is ultimately my

⁵ In this I follow Bratman (1992), p. 9, against Stalnaker (1984).

⁶ Thanks to Eleonora Orlando and Agustín Rayo for advice concerning the way to express this point.

taking *h* to be true which explains what I do). The same goes for belief-contravening assumptions for the sake of the argument – the type of assumptions one finds in *reductio* arguments. Again, as it has often been emphasized in the belief revision literature, the starting point of a *reductio* should never be taken to reflect our actual epistemic state.⁷ Indeed, given that contexts containing the same epistemic certainties are taken to be identical, it would be a mistake to think that my engaging in a *reductio* counts as the type of circumstance that by itself defines a new, different context. It does not. 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. All it means is that I have decided not to disclose my real convictions, out of prudence; in order to make this decision I was bound to take those very same convictions into account. Hence the circumstances in which I utter such words do not constitute a different context, in the sense defined here.

- (3) Related to the previous point, ‘context’ should not be understood in the way so-called contextualism in epistemology does.⁸ According to the contextualist, contexts are individuated by reflecting on the knowledge ascriptions that come out true or false in each case, but this is not what the current project is about. Moreover, here I shall not be concerned with the concept of knowledge at all – at least not in the way this concept is used in mainstream epistemology; more generally, contexts in my sense are not to be thought of as involving the subject’s assessment of potential ascriptions of various epistemic attitudes (such as acceptance or belief) to *different* agents. If anything, my understanding of contexts is closer (though not identical) to what Hawthorne (2004) or Stanley (2005) call ‘practical environments’ – of the primary epistemic subject, rather than of those who make subsequent epistemic attributions.⁹ Related to this, in this paper I will not be concerned with iterated epistemic attitudes, or with the concomitant discussion about issues such as transparency or luminosity (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 should not be hard to extend the present account in order to take care of these topics.

⁷ Here I follow the standard treatment of *reductio* arguments in the belief revision literature (see for instance Levi (2004), ch.1). *Reductio* arguments require that we consider the *contraction* of an epistemic state for the sake of the argument – which has to be carefully distinguished from an actual, *bona fide* contraction. In this paper I am just extending this idea to the possibility that *reductio* procedures (with their concomitant revisions for the sake of the argument) be also contextual. See footnote 18 for further precisions on how to assess the relation between counterfactual conditional claims and ‘contextual’ conditionals.

⁸ Cf. Cohen (2000), or DeRose (2002), among others.

⁹ In the light of this, it can be argued that it might be more convenient to use the labels ‘environmental acceptance’ and ‘environmental change’, rather than ‘contextual acceptance’ and ‘contextual changes’ – but I suspect these neologisms might bring about further confusions of their own.

- (4) 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? Similarly, what if for every contingent potential assumption an agent can conceive of, there is a particular context in which it is absent? In any of these cases, bona fide beliefs would be reduced to beliefs about logical truths.¹⁰

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. 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 – some such ‘doubts’ are simply not compelling to most of us. Speaking by myself, I am ready to acknowledge that, at least under some circumstances, 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 any of 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.

- (5) 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 under different circumstances, 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 models; in particular, I agree that epistemic rationality imposes the ideal of consistency and deductive closure on us, in the sense that we do long for epistemic states that exhibit such features, even if they are unattainable in practice. But insofar as the idea of rationality and rational agency does not force us to say that we should have not doubts whatsoever,¹¹ 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 closure should be relativized to contexts.

In the next section I shall present the basic structure of the model I favor, which seeks to capture, at a representation level, the main intuitions delineated so far. Let me emphasize that I shall not attempt to develop a comprehensive, complete theory about the agent’s real

¹⁰ Throughout this paper I shall assume that classical logic holds. See point (5) for some observations on deductive closure.

¹¹ See an interesting discussion of this point in Christensen (2004), chapter 6.

epistemic state; in particular, I shall bypass what we might dub ‘the ontological question’ on beliefs and acceptances, as mentioned on p. 5: the analytic tool I offer here will be compatible with many different approaches on the exact nature 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, as we shall see, I shall propose to model epistemic attitudes by means of sentences of a representation language L , such that the sentences of L (at the representation level) *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 (such as acceptance, rejection or suspension of judgment), or an analysis of her semantic assumptions regarding such sentences, will not be a goal in itself.

Let me also add a few reflections on the purposes of the formalism. In addition to enhancing our understanding of the way beliefs and acceptances interact in particular agents (possibly ourselves), by offering a formal reconstruction of epistemic states we may gain awareness about the rational way to proceed in the future. Thus, we shall see that, by specifying the relevant parameters, it may become clearer whether an agent should or should not accept a given hypothesis in a context, by her own lights, and it may also become apparent whether a split of contexts should take place – which may in turn be consequential at the time of making further epistemic decisions.

5. Modeling beliefs and acceptances

Let me 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 then be modeled by a (non-empty) convex set $\Delta_{X,t}$ of credal probability functions P_k over the sentences of L . (I shall get rid of sub-indices when there is no risk of confusion). Let \mathbf{T} be the set of all theories T_i of L , and define $f: \Delta \rightarrow \mathbf{T}$ such that $f(P_k) = \{\alpha \in L: P_k(\alpha) = 1\}$.¹² 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 $T_i \in \mathbf{T}$, $T_i = A_i$ 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 A_i , or the set of acceptances of the agent in context i . By extension, we shall also say that Δ_i is correlated with A_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.

Next, define $K = \bigcap_j A_j$, for all contexts 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

¹² Thanks to an anonymous referee for the *Formal Epistemology Workshop 2008* for suggesting this formulation.

might be further acceptance sets (in addition to K) embedded in others, though this is not mandatory.

In addition, define $M = \cup_j A_j$, for all contexts 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 any case, this fact does not enable agents to deduce explicit contradictions, because M need not be deductively closed – hence, the problematic sentences (in case there are any) can remain isolated from one another, so to speak.¹³ 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; 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 agents can be 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.¹⁴ Clearly, contextual and full doubts are probabilistically ordered.

The model trades on well-known attempts to refine standard Bayesianism.¹⁵ Within such refinements, 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. (I will address possible problems with the use of intervals at the time of updating probabilities in due course). In addition, we obtain a straightforward distinction between different senses of probability 1, insofar as contextual acceptances receive probability 1 without thereby being full beliefs.¹⁶ In this way, we highlight the idea

¹³ Still, we might also demand that agents who find themselves holding inconsistent sets of acceptances at t seek to eliminate the conflicting assumptions through one or more contextual contractions; I shall address the topic of contractions very briefly in section 8.

¹⁴ Thanks to Paul Pedersen for pointing out a problem with a previous version of this definition.

¹⁵ 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). I address some problems associated with the updating of sets of probabilities at the end of section 6.

¹⁶ 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 inner core at

that there is a peculiar continuity between an agent’s state of certainty and doubt, as we wanted. 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.

Let me stress that sentences of L (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 basic epistemic state, and hence of a derivative nature.¹⁷ This qualification should suffice to block the following potential concern. Suppose 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) *gets 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 i 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

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)).

¹⁷ 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.

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.¹⁸

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 . Related to this, the present framework should not be meant to imply that probability assessments have *ontological priority* over certainties; after all, we could have arrived at the very same structure by means of a different strategy. To wit, we could have well begun by presenting a bunch of acceptance sets A_j , which would then be supplemented by corresponding sets of probability measures. The favored approach, however, takes probability functions to be more basic *at the representation level*, but only for pragmatic reasons¹⁹; as I can see it, by proceeding thus we obtain a simpler (and, to some extent, more elegant) epistemic model.²⁰

6. Voluntary epistemic changes: The case of contextual expansions

In the remaining sections of this paper I shall explore how the model behaves at the time of representing epistemic changes. 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

¹⁸ A more complete treatment of this potential 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 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.

¹⁹ See the cautionary note about the description of the agent's real epistemic state, at the end of section 4.

²⁰ Incidentally, notice that within the present framework probability 0 and 1 collapses with (epistemic) impossibility and necessity, respectively (in each context). In other words, here we cannot allow that a sentence of L represents a serious possibility without receiving positive probability by 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 16.

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 . The present section will be devoted to contextual expansions.

In previous work 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.²¹ In other words, I have suggested that we conceive of IBE as a decision theoretic exercise, in which we focus on the epistemic gain we are able to obtain. This requires our paying attention to both personal probabilities and epistemic utility functions, where epistemic utilities, in turn, can be assumed to incorporate features of hypotheses such as simplicity, unification power, fertility, accuracy, or predictive force, among others. 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 incorporates references to several kinds of contextual indices, such as caution thresholds, or contextual weights for the several dimensions that compose an epistemic utility function. Yet in addition to all such indices, if our discussion from sections 3 and 4 was on the right track, different research processes may take different sets of epistemic certainties for granted. 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 propositions) 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 Δ_b correlated with context b (for instance, recall the moment Jill Jones sought to assess which explanation for fetomaternal tolerance was the *best*, given her epistemic state at the time). 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 brand of cognitive decision theory we adopt. Before building the set of possible options the agent might feel compelled to seek for new evidence, which might lead her to update all measures in Δ_b through Bayesian conditionalization. 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 options as determined by the decision theoretic problem), *for each probability measure 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 1.²² Notice that, depending on the case, the

²¹ Cf. Cresto (2008). Well known approach to cognitive decision theory are found in Levi (1980), Maher (1993), or van Fraassen (1989, 2002).

²² It might be contended that allowing probability functions to be modified in the way I am advocating here commits 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

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.²³

In what follows I shall not be concerned with the development of an epistemic utility function, 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. Actually, all we have to assume is that, regardless of the details, there is *some* account we can use at the time of deciding whether a given probability function licenses the acceptance of a particular hypothesis. Rather than focusing on what makes a proposition *acceptable*, let me focus then on how to proceed once we agree that our coming to accept it is legitimate – in short, let me concentrate on how to update acceptance sets and contexts (and, eventually – as we shall see in section 10 – full beliefs). Thus, notice that although I favor a decision theoretic account for acceptance, the proposal that follows is neutral concerning the mechanism by which epistemic decisions are made.

Let $\Delta_{X,t}$ represent agent X 's epistemic state at t , as usual; also, for any i , let Δ_i be correlated with T_i in the usual manner; assume $\Delta_b \subseteq \Delta_{X,t}$ is nonempty.²⁴ We 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

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.

²³ 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 so as to give probability 1 to hypothesis h iff h 's expected epistemic utility is highest for *all* such members (thanks 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.

²⁴ 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.

α in a particular context a (for α consistent with T_b) if and only if the following conditions are met:²⁵

(1) Closure

We want the result of the expansion to be another epistemic state:

$\Delta'_{x,t}$ is a (nonempty) convex set of probability measures.

(2) Success

We want some function in Δ_b to be updated so that now it picks out $T_a = Cn(T_b \cup \{\alpha\})$, where ‘ Cn ’ is, as usual, the Tarskian operator of logical consequence. Formally:

There is set $\Delta_a \subseteq \Delta'_{x,t}$ such that, for any sentence γ :

(a) There exists some P_k in Δ_a with $P_k(\gamma) = P_m(\gamma/\alpha)$, for some P_m in Δ_b ;

(b) For every other P_n in Δ_a : either

i. $P_n(\gamma)$ is a mixture of $P_m(\gamma/\alpha)$ (for $P_m \in \Delta_b$) and the values assigned to γ by other functions in $\Delta_{x,t}$ that pick out proper supersets of $Cn(T_b \cup \{\alpha\})$ (in case such functions exist); or

ii. $P_n(\gamma) \in \Delta_s \subseteq \Delta_{x,t}$ for Δ_s correlated with $T_s = Cn(T_b \cup \{\alpha\})$; or

iii. P_n is a mixture of functions satisfying (2.a), (2.b.i), or (2.b.ii).

(3) Non-vacuity and merging

We want some real change to occur: if α was already accepted in context b , the definition should not be satisfied. On the other hand, if α was not accepted in b but $Cn(T_b \cup \{\alpha\})$ was already an acceptance set at t , a merging of contexts should be obtained. Formally:

If there exists already $\Delta_s \subseteq \Delta_{x,t}$ for Δ_s correlated with $T_s = Cn(T_b \cup \{\alpha\})$, then $\Delta_b \neq \Delta_s \subseteq \Delta_a$.

(4) Conservativeness

We do not want unnecessary losses of former probability functions in the new epistemic state:

For all P_i in $\Delta_{x,t}$ and all sentences γ :

$P_i \notin \Delta'_{x,t}$ iff both $P_i \in \Delta_b$ and $P_n(\gamma) = P_i(\gamma/\alpha)$, for some P_n in Δ_a .

(5) Minimality

We want the change from $\Delta_{x,t}$ to $\Delta'_{x,t}$ to be the smallest shift that fulfills (1)-(4):

For all P_j in $\Delta'_{x,t}$:

(a) Either P_j was in $\Delta_{x,t}$; or

(b) $P_j \in \Delta_a$; or

(c) P_j is a mixture of elements satisfying (5.a) or (5.b).

(See the next section for an illustration of how these conditions work in a concrete example.)

²⁵ Thanks to Agustín Rayo for useful advice concerning the presentation of these conditions.

Let us examine some consequences of this definition. First, it is straightforward from (2) that $T_a = Cn(T_b \cup \{\alpha\})$, as desired. In other words, for all P_n in Δ_a and any sentence γ , $P_n(\gamma) = 1$ iff $\gamma \in Cn(T_b \cup \{\alpha\})$.

(Proof: If $P_n \in \Delta_s$ (as defined in (2.b.ii)) the result is trivial. Suppose $P_n \notin \Delta_s$, and suppose also that P_n is as in (2.a). If $\gamma \in Cn(T_b \cup \{\alpha\})$, then $P_n(\gamma/T_b \ \& \ \alpha) = 1$. But $P_n(T_b) = P_m(T_b/\alpha) = P_n(\alpha) = P_m(\alpha/\alpha) = 1$, for some P_m in Δ_b . Hence $P_n(\gamma/T_b \ \& \ \alpha) = P_n(\gamma) = 1$. By a symmetric reasoning, if $\gamma \notin Cn(T_b \cup \{\alpha\})$, then $1 \neq P_n(\gamma/T_b \ \& \ \alpha) = P_n(\gamma)$. Now we have only two cases left to consider. Suppose P_n is as in (2.b.i) or (2.b.iii). Then P_n is in the convex hull generated by the functions already examined and, possibly, measures that pick out proper supersets of $Cn(T_b \cup \{\alpha\})$. Clearly, any mixture between measures Q and Q' will give probability 1 to γ iff γ receives probability 1 from both Q and Q' , hence any mixture in Δ_a gives probability 1 to γ iff γ belongs to $Cn(T_b \cup \{\alpha\})$. ■)

Second, (2.a) and (3) entail that at least some function in Δ_b has been updated. This guarantees that we do not use the expression “ X comes to accept α while reasoning on the basis of T_b ” just because α was already in T_b . By (2.a), if the inclusion in (3) is strict so that $\Delta_a = \Delta_s$, then at least some element in Δ_s should be identical to the updating of a function in Δ_b . (3) also guarantees that, had there been a superset of T_b containing α (but no other additional sentence) in $M_{X,t}$, a fusion of contexts would be obtained. In short, in order to be fulfilled this definition requires that some real probability change occur, regardless of whether α had already been in some acceptance set before t' . Notice that 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.

On the other hand, α cannot become an element of acceptance sets *properly included* in T_b ; functions that pick out proper subsets of T_b are not in Δ_b , and hence will not be shifted. 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 .

In any case, notice that our definition does not require the modification of *all* measures in Δ_b . By (4), elements of Δ_b that remain unchanged are not lost. 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, and assuming there was no context associated with $Cn(T_b \cup \{\alpha\})$ at t , 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 $(\mathbf{T}_M \cup \{Cn(T_b \cup \{\alpha\})\}) \setminus \{T_b\}$. On the other hand, if $Cn(T_b \cup \{\alpha\})$ was already an acceptance set at t , we will have either $\mathbf{T}'_M = \mathbf{T}_M$ or $\mathbf{T}'_M = \mathbf{T}_M \setminus \{T_b\}$, depending on the details. This is as it should be, if we want to allow for the possibility that the very space of contexts change.

In the next section I shall go back to Jill’s case, from section 2, to illustrate in a more detailed fashion how the proposal is meant to work. But first, let me address a potential worry concerning the update of convex sets of probabilities. Note that if the set of

probability measures correlated with a particular context is such that the probability of q in the context adopts all values in $[x,y]$, and there is a finite partition $p_1\dots p_n$ with positive priors, such that q given p_i adopts all values in $[x',y']$, for $x' < x \leq y < y'$, then conditionalizing on p_i ($i=1\dots n$) will have the effect of enlarging the set of q 's values in the context. This phenomenon is known as (strict) *dilation*.²⁶ In the literature, $p_1\dots p_n$ are typically taken to be observational statements, but given the acceptance mechanism discussed above, we should also consider the possibility that they stand for theoretical hypotheses. Reactions to dilation vary widely; while some authors do not seem to be particularly bothered by it, others consider it disastrous for systems that attempt to model epistemic states by means of so-called 'imprecise' probabilities.²⁷

I do not have a knockdown argument to the effect that the advantages of the use of convex sets of probabilities outweigh the difficulties.²⁸ Still, there are a number of considerations that suggest that the negative impact of the phenomenon is not as great as one might think at first blush. Notice, first, that even though some particular examples are indeed hard to swallow, dilation is not always intuitively anomalous. Sometimes the logical ties between q and p_i make the result intuitively acceptable.²⁹ At other times, the agent's unnatural ignorance of the (seemingly obvious) probabilistic independence between q and p_i makes the result, once again, intuitively as it should be. (Consider for instance van Fraassen's (2006) description of a student who allegedly ignores whether there is probabilistic dependence between meteorological phenomena and his own performance in school tests; it could be contended that typical agents are not as ignorant as to ensure that this type of example generalizes – and if they are, then perhaps they deserve what they get). More importantly, it is far from obvious that the problem will persist in the long run.³⁰ Even acknowledging that classical theorems on the asymptotic merging of Bayesian posterior probabilities do not apply to the present framework (given that, among other things, we are dealing with uncountably many priors)³¹ we can still seek to ensure that the gathering of new data fulfills further conditions that guarantee that no asymptotic dilation occurs – at least for cases of dilation related to the acquisition of new empirical evidence (on this cf. especially Herron, Seidenfeld and Wasserman 1994, 1997).

²⁶ Thank you to an anonymous referee for pressing this problem. Cf. for instance (Walley, 1991, p. 299); (Seidenfeld and Wasserman, 1993); (Herron, Seidenfeld and Wasserman, 1994, 1997); (van Fraassen, 2005, 2006); or (White, 2008), among others; cf. (Sturgeon, 2008) for an attempt to deal with (White, 2008).

²⁷ For example, the existence of dilation leads van Fraassen (2006) to advocate for the need to constrain opinions by means of 'hidden variables' that would ensure the stability of prior graduate beliefs, whereas Elga (mn) and White (2008) take dilation to support the claim that probabilities should be sharp.

²⁸ Cf. footnote 15 for a brief comment on some of the advantages.

²⁹ For example, for any r and p_i probabilistically independent such that the probability of p_i is 0.5 and the probability of r adopts all values in $[0,1]$, we also have $q=(r \& p_i) \vee (\neg r \& \neg p_i)$ with sharp probability 0.5. If we conditionalize by p_i , however, the posterior probability of q adopts all values in $[0,1]$. Cf. Herron, Seidenfeld and Wasserman (1994).

³⁰ Thanks to Horacio Arló Costa for pointing to me this line of response.

³¹ Cf. (Savage, 1954), or (Blackwell and Dubins, 1962).

7. Jill’s example revisited

Let us go back for a moment to Jill Jones, our biologist from section 2. Consider 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 the 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 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} :

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

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

$$\Delta_k = \{P_i: P_i \text{ is a mixture between } P_1 \text{ and } P_j, \text{ for all } P_j \in \Delta_b\}$$

(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 $t \vee u$.³²

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 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 her party tomorrow – which results, once again, in s ’s having a

³² Recall that, within the present model, different possible circumstances are rendered as constituting the same epistemic context if they lead the agent to embrace the same set of epistemic assumptions. Cf. section 4.

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.

Now suppose Jill wonders whether her current epistemic state makes it reasonable for her to adopt the best explanatory hypothesis for the problem of fetomaternal tolerance. In agreement with the details given in section 2, she reasons within context b , and considers whether to expand by t , by u , or not expand at all, as the case may be. 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 to them epistemic utilities of 0.4 and 0.8, 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.³³ 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

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

(Of course, the situation would have been rather different if P_2 and P_3 did not agree on the values they assigned to sentences 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:

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

(hence $P_2' = P_2$); 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_1 Jill holds Δ'_{J,t_1} with contexts a' , b' , b'' and k' , as follows:

$$\Delta_{a'} = \Delta_a = \{P_1\}$$

$$\Delta_{b'} = \{P_3'\}$$

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

$$\Delta_{k'} = \text{All remaining mixtures}$$

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

³³ 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 $eu(\sim h)$ is 0.

herself while doing research (as far as she is concerned), as we wanted. The structure of $\Delta'_{J,t}$ 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.

8. Contextual contractions

Contextual *contractions* proceed along similar lines, although, as is natural, in this case there are several additional problems to take care of. A detailed analysis of contractions would exceed the limits of this paper, so it will be left for future work; here I shall only offer a partial elucidation of this notion, by means of a partial set of necessary conditions. 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; I shall not examine other mechanisms that may lead to the decision to remove some sentence α from a specific theory.

We say that a change from $\Delta_{X,t}$ to $\Delta'_{X,t'}$ is a contextual contraction by α for a particular context (by which we mean to say that X ceases to accept α at t' , in some context) only if:

(1) Closure

$\Delta'_{X,t'}$ is a (non-empty) convex set of probability measures.

(2) Non-vacuity

There is some set $\Delta_a \subseteq \Delta_{X,t}$ such that $\alpha \in T_a$ and $\Delta_a \not\subseteq \Delta'_{X,t'}$

(3) Success

There is $\Delta_b \subseteq \Delta'_{X,t'}$ such that:

- a. For every P_i in Δ_b , $P_i(\alpha)$ is in some interval $(x,y) \subseteq (0,1)$.
- b. For every P_i in Δ_b and every sentence γ , if γ is either logically independent of α or probabilistically independent of α relative to P_i , then 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 , there is some P_i in Δ_b such that, if γ is either logically independent of α or probabilistically independent of α relative to P_i , then $P_i(\gamma) = P_j(\gamma)$.

(4) Merging

If there exists already $\Delta_s \subseteq \Delta_{X,t}$ for Δ_s correlated with $T_s = T_b$, then $\Delta_s \subseteq \Delta_b$.

(2) guarantees that some real change has occurred from t to t' , while (3.a) guarantees that α is no longer accepted in the new context; as usual, if α is tautological, condition (3.a) cannot be satisfied. (3.b) makes sure that there is no arbitrary new information in Δ_b , whereas (3.c) guarantees that there are no unnecessary losses. Of course, this does not help us to uniquely determine the new set $\Delta'_{X,t'}$. Different approaches could be adopted here,

which would lead to different instructions as to how to fix Δ_b (and hence $\Delta'_{X,t}$). Notice that, so stated, (3) allows that the same function in Δ_a yield multiple ‘daughters’ in Δ_b , for different values of α . As with expansions, by (4), if Δ_b ends up being correlated with a theory that was already in $M_{X,t}$, we will have a fusion of contexts.

9. 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 hypothesis 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.

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 sentence 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 expansions: the reason is that each contextual change does not have *the goal* of modifying her set of beliefs.³⁴

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 Δ^m_{X,t_m} to Δ^n_{X,t_n} such that:

- 1) For some P_i in Δ^m_{X,t_m} , $P_i(\alpha) \neq 1$.
- 2) For every P_j in Δ^m_{X,t_m} and all γ , there is some P_k in Δ^n_{X,t_n} such that $P_k(\gamma) = P_j(\gamma/\alpha)$.³⁵

³⁴ 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 .

³⁵ Note that different measures in Δ^m_{X,t_m} may well yield the same function in Δ^n_{X,t_n} ; moreover, if α was already accepted in some contexts, some measures will not be shifted.

3) There are no other functions in $\Delta^n_{x,t}$.

Or:

(B) There is a sequence of consciously implemented contextual expansions $\Delta^1_{x,t}, \dots, \Delta^m_{x,t}, \Delta^n_{x,t}$ (as defined in section 6), such that:

- a. For some P_i in $\Delta^1_{x,t}$, $P_i(\alpha) \neq 1$; and
- b. For all P_j in $\Delta^n_{x,t}$, $P_j(\alpha) = 1$.

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 $\Delta^m_{x,t}$ to $\Delta^n_{x,t}$ is justified in the manner explained in section 6, 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.³⁶ 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 agent has not consciously chosen to implement, and which have the effect of erasing a given sentence from every acceptance set in M .³⁷

10. Conclusions

Let me summarize briefly what we have achieved. In section 2 I have offered a case study, which suggested the convenience of embracing 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 agreement with this, in section 4 I have presented the intuitive distinction between contextual epistemic assumptions and epistemic assumptions held in all contexts; I have 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 have 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.

³⁶ 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. footnote 34.

³⁷ It could be complained that, typically, the belief/acceptance distinction, as used in the literature, is meant to highlight the difference between degree-like states, which are not voluntary, and voluntary states, which are representable non probabilistically. Hence we may wonder how we can take acceptances to be voluntary, given that they are described as probabilistic states. (Thanks to an anonymous referee for raising this issue). To answer this objection we should recall that what is actually voluntary is the update of a (set of) probability measures to 1, in a particular context (i.e., our *coming to accept*), or, symmetrically, the update of a (set of) probability measures with probability 1 to less than 1 (our *ceasing to accept* – or *coming to doubt*). Moreover, notice that by coming to accept a particular sentence what the agent voluntarily adopts is a new (contextual) *certainty*; certainties are indeed “probabilistically representable,” but they are of course given probability 1. Thus, acceptances are degree-like states only in a degenerate sense, as they can only receive probability 1.

In sections 5 to 9 I developed the bulk of my formal proposal. The model aims to capture some phenomena that are intuitively clear at the level of the agent's *real* epistemic state, while being neutral with regard to the ontological problem of beliefs and acceptances. I have 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 could be situated – circumstances that add to the constitution of what the agent takes to be some relevant context. I have 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 may expand M (the set of sentences accepted at *some* context) and not 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 9 I have pointed out that the model allows for spontaneous changes in K . In addition, I have 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 sentence 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, and it makes room for the idea that our voluntary epistemic life proceeds in a 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 well known difficulties 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.

Acknowledgments

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 *Fifth Principia International Symposium* (Florianopolis, August 2007), at a discussion session at the *GAF (Grupo de Acción Filosófica)*, Buenos Aires, September 2007), and at the *Formal Epistemology Workshop 2008* (Madison, Wisconsin, May 2008; with comments by Bernhard Nickel). I am grateful to the participants and commentator for helpful discussion and suggestions.

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.
- Blackwell, D. and Dubins, L. (1962). “Merging of opinions with increasing information”. *Annals of Mathematical Statistics*, 9:235-244
- 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.
- Bratman, M. (1992). “Practical reasoning and acceptance in a context”. *Mind*, 101: 1-15
- Christensen, D. (2004). *Putting Logic in its Place. Formal Constraints on Rational Belief*. Oxford: Clarendon Press.
- 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.
- 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.
- 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.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.

- Hawthorne, J. and Stanley, J. (2008). “Knowledge and Action”. *The Journal of Philosophy*, 105 (10).
- Herron, T, Seidenfeld, T., and Wasserman, L. (1997). “Divisive Conditioning: Further Results on Dilation”. *Philosophy of Science*, 64: 411-444.
- Herron, T, Seidenfeld, T., and Wasserman, L. (1994). “The Extent of Dilation of Sets of Probabilities and the Asymptotics of Robust Bayesian Inference”. *Proceedings of the Philosophy of Science Association* 1994, Vol.1, pp. 250-259.
- 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.
- Savage, L.J. (1954). *The Foundations of Statistics*. New York: Wiley.
- Seidenfeld, T. and Wasserman, L. (1993). “Dilation for Sets of Probabilities”. *Annals of Statistics*, 21: 1139-1154.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, Mass.: The MIT Press.
- Sturgeon, S. (2008b). “Confidence and Coarse-Grained Attitudes.” Manuscript. Available at http://www.fitelson.org/few/few_08/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.
- Van Fraassen, B. (2005). “Conditionalizing on Violated Bell’s Inequalities”. *Analysis*, 65: 27-32.
- Van Fraassen, B. (2006). “Vague Expectation Value Loss”. *Philosophical Studies*, 127: 483-491.
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. London: Chapman and Hall.
- White, R. (2008). “Evidential Symmetry and Mushy Credence”. Manuscript. Available at http://www.fitelson.org/few/few_08/white.pdf.