



BIROn - Birkbeck Institutional Research Online

Vorms, Marion and Hahn, Ulrike (2019) In the space of reasonable doubt. Synthese , ISSN 0039-7857. (In Press)

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/30538/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html>
contact lib-eprints@bbk.ac.uk.

or alternatively

IN THE SPACE OF REASONABLE DOUBT

Abstract

This paper explores ‘reasonable doubt’ as an enlightening notion to think of reasoning and decision-making generally, beyond the judicial domain. The paper starts from a decision-theoretic understanding of the notion, whereby it can be defined in terms of degrees of belief and a probabilistic confirmation threshold for action. It then highlights some of the limits of this notion, and proposes a richer analysis of epistemic states and reasoning through the lens of ‘reasonable doubt’, which in turn is likely to supplement the DT framework. The strategy consists in fighting on two fronts: with DT, the paper claims that there is no absolute (*i.e.* decision-independent) notion of ‘reasonable doubt’ but, *pace* DT, it shows that reasonable doubt cannot be accounted for only in terms of degrees of belief and probabilistic threshold. We argue that the lens of reasonable doubt sheds light on aspects of belief dynamics, as well as of the nature of epistemic attitudes, which are often obscured by belief-centred approaches. In particular, when it comes to acknowledging the necessary ignorance and irreducible uncertainty that we face in our everyday-life decisions, studying the various facets of doubt rather than focusing on what can be believed, enables one to do justice to the richness and diversity of the mental states in play.

Introduction

When is doubt reasonable? Doubting is widely acknowledged as epistemically virtuous, and as shielding us from dogmatism and fanaticism. In that sense, one could be tempted to argue that doubt is *always* reasonable — after all, no empirical hypothesis is ever 100% certain, hence an uncompromising truth-seeker should consider the falsity of any

of them as a possibility. However, in most real-life situations, waiting until complete certainty is achieved before taking action seems an unreasonable and impractical attitude. The issue of what makes doubt (un)reasonable therefore seems to arise primarily in the practical domain, where some decision has to be made that depends on one's conclusions regarding the hypothesis at stake. Typically, should one convict or acquit a person accused of having committed a given crime? But also, should one vaccinate a child? Take steps to mitigate climate change? Or should one rather doubt the hypotheses grounding the arguments in favour of such actions?

Even when no practical decision appears to be immediately at stake, following the sceptic's recommendation of keeping one's mind open to any alternative hypothesis to the most widely held and/or best supported by the evidence, need not always be reasonable. Open-mindedness and doubt are often misused as slogans by "merchants of doubt" (Oreskes & Conway 2010), who question the authority of institutional science in bad faith, as well as by the advocates of those highly unreasonable conjectures that fall under the fuzzy category of "conspiracy theories". Under the guise of critical thinking and anti-dogmatist approaches, one finds "free minds" invited to doubt the official theory (according to which, for instance, the Earth is round, or the 9/11 attacks were plotted by Al-Qaeda) — if not to adopt a specific alternative theory (according to which the Earth is flat or the 9/11 attacks were plotted by the Mossad and the CIA) ¹. Where does the boundary lie, between epistemic over-cautiousness verging on social paranoia, and the fair stubbornness of the juror embodied by Henry Fonda in *Twelve Angry Men*², who refuses to convict the defendant without first questioning in every detail the prosecution's story (and especially the coherence and credibility of all pieces of supposedly incriminating evidence), even though this story is prima facie highly plausible? In both cases, one rejects, at least temporarily, the official version (the most widely held, authorised, consensual, and probably the best supported by the available evidence); one highlights its flaws, and explores any conceivable alternative, even

¹ As Stephen John emphasised in a personal communication, it is very important to keep the issues of rejecting a theory and of accepting another separate, since it is possible to undermine, *e.g.*, climate science without providing an alternative positive theory. Many conspiracy theories thus rest on an illegitimate assumption that the burden of proof is on the side of the official theory. See Hahn & Oaksford 2007 for reflection on the use of the notion of burden of proof in argumentation beyond the legal domain.

² Movie by Sidney Lumet (1957), from a TV play by Reginald Rose (1954).

though prima facie implausible, scenario. How are we to account for the ‘common-sense’ intuition, according to which what seems delusional in one case is reasonable in the other?

An obvious candidate seems to be a decision-theoretic: A natural answer to the question above consists in emphasizing the importance of the decisional context. In the case of the juror, a man’s life being at stake, it can be reasonable to keep doubting his guilt even in the presence of strongly incriminating evidence, while such obstinacy may become unreasonable in situations where the stakes are lower (or where nothing practical seems to be immediately at stake). Representing the degree of confirmation of the hypothesis in question in probabilistic terms, one can thus define ‘reasonable doubt’ in terms of a probabilistic confirmation threshold, set according to the context and the associated utilities.

Being thus closely tied to rational action, what can reasonable *doubt* bring to the analysis of reasonable *belief*? The notion of belief is central to at least two domains of contemporary epistemology: first, the study of the dynamics of belief (or rational belief change), and, second, the branch of mainstream epistemology concerned with epistemic attitudes and their justification.³ Those two domains rarely communicate, and the ‘beliefs’ they posit as the primitive elements of mental life have incompatible definitions — despite some attempts at conciliating them.⁴ The two domains have in common, though, that they have overlooked the notion of doubt.⁵ One of our goals in this paper is to show that the lens of reasonable doubt sheds light on aspects of belief dynamics, as well as of the nature of epistemic attitudes, which are often obscured by belief-centred approaches.

³ See for instance Bonjour & Sosa 2003, Dancy 1985, Pollock 1986, Pritchard 2005.

⁴ See footnote 46 below.

⁵ One contemporary exception is Thagard (2004), but this paper remains marginal, and his definition of doubt is all in all quite cursory. To our knowledge, the most in-depth recent analysis of doubt is Salmon’s (1995), but his proposal makes sense against the background of issues in the philosophy of language and logic, rather than in an epistemological and psychological framework as ours. Jane Friedman’s (2017, forthcoming) works linking inquiry with suspension of judgement notably put an emphasis on the importance of those epistemic states and actions that are generally ignored by belief-centred approaches. Finally, it is worth mentioning that, while the present paper was under review, two papers on doubt have appeared (Lee 2018; Moon 2018). The conceptions of doubt that they defend are different from ours, but the issues they address are mostly orthogonal to the main focus of the present paper, which places discussion of their views beyond its scope.

Our strategy will consist in fighting on two fronts: with the decision-theoretic view (from now on, “DT”), we claim that there is no absolute (*i.e.*, decision-independent) notion of ‘reasonable doubt’ but, *pace* DT, we show that reasonable doubt cannot be accounted for only in terms of degrees of belief and probabilistic threshold. In particular, when it comes to acknowledging the necessary ignorance and irreducible uncertainty that we face in our everyday-life decisions, studying the various facets of doubt rather than focusing on what can be believed enables one to do justice to the richness and diversity of the mental states in play. As we will see, depending on the context and the hypotheses at stake, doubting might indeed imply substantially diverse epistemic attitudes and actions.

Emphasizing such diversity, and the complexity of the components of the (un)reasonableness of doubt, should contribute to highlight the similarities between cases of applied scepticism as diverse as the juror’s and the conspiracy theorist’s ones. In particular, this will imply questioning the distinction between cases with practical stakes, from ‘purely’ theoretical ones.

In what follows, we adopt the decision-theoretic framework. From this starting point, we examine several intuitions underlying the notion of reasonable doubt. Rather than systematically questioning its power and relevance, our strategy throughout the paper consists in conceding as much as possible to the DT framework. This will lead us to finally highlight some of its limits, and to propose a richer analysis of epistemic states and reasoning through the lens of ‘reasonable doubt’, which in turn is likely to supplement the DT framework. In the first section, we sketch the minimal, *prima facie*, definition of reasonable doubt that can be given in DT terms. In section 2, we deepen our analysis of the notion of doubt as it arises in this framework. Section 3 takes a detour through an analysis of ‘full belief’, examining to what extent it escapes the DT framework (as being apparently independent from any decisional context), and section 4 explores whether the DT framework can account for all aspects of the (un)reasonableness of doubt — especially when no practical decision appears immediately at stake. Although our conclusions at this point will not go frontally against the DT framework, those analyses will finally lead us to highlight some intrinsic limits of the consequence-dependent probabilistic threshold view of reasonable doubt in section 5.

1. The decision-theoretic (DT) image of reasonable doubt

In everyday life, as well as in the scientific, legal, medical, or policy-making contexts, the informational basis for making decisions is often complex, heterogeneous, incomplete, and partially contradictory. In most situations, practical necessity compels us to take action upon partially confirmed hypotheses: although complete certainty hasn't been reached regarding the relevant hypothesis H (and is most of the time a long way off), one must act based thereupon. Let us call such action based upon H , *acceptance of H* .

1.1 *Acceptance as context-dependent*

Before spending a day out without knowing whether it will rain or not, one should for instance decide whether to accept that it will rain (act upon this hypothesis by carrying one's umbrella) or not (leave one's umbrella home). As is obvious from this example, the degree to which a hypothesis is supported by the evidence isn't enough to determine whether one should accept it: both the available evidence (provided by the weather forecast and other relevant sources), and the decisional context including the utilities associated to the expected consequences of the various possible actions, have to be taken into account (factoring in, *e.g.*, the relative waterproofness of one's coat, and the weight and cumbersomeness of one's umbrella). Similarly, a doctor deciding whether to prescribe a medication with potentially severe side-effects (assuming such medication being the only chance to cure a serious disease that her patient shows some symptoms of), or a policy-maker deciding whether to implement a costly policy aimed at shortly reducing greenhouse gases emission, should in principle make a cost-benefit analysis based on the available information, and decide on which hypothesis to act (that the patient actually suffers from that disease, that the experts' forecast, according to which the current emission level of greenhouse gases would have irreversible effects if it does not fall in the short term, is accurate). One can thus assume that there is a confirmation threshold, above which one should accept a given hypothesis — *i.e.*, take action based on it. Below such threshold, one should not accept hypothesis H .

Depending on the situation, this may imply rejecting H (and accepting $non-H$ ⁶), or rather seeking more information before accepting or rejecting H .

Note that the threshold for acceptance is context-dependent, meaning that the same hypothesis might be acceptable *vis-à-vis* one decision, while it should be rejected *vis-à-vis* another one. Consider the hypothesis H according to which it will rain this afternoon. If the decision at stake concerns whether or not to carry one's umbrella, assuming that the umbrella is light and one's coat not waterproof, it may seem reasonable, even with a low $p(H)$, to accept H — to carry one's umbrella, acting as if it was going to rain. But if the decision concerns whether to take one's plants out on the rooftop or not (assuming this is a heavy task and one will have to bring them back in at night), then one might not accept H , and wait until a day when the likelihood of rain be higher. The confirmation degree of H is the same in the two situations, but the threshold for acceptance moves.⁷

What does 'reasonable doubt' mean, in such a framework? One straightforward answer would be to define it just as the complementary to rational action: there is reasonable doubt as to H when the acceptance threshold for H hasn't been met. *Prima facie*, this is all the DT framework has to say about 'reasonable doubt'. Let us now consider the task of jurors in criminal trials from this perspective.

1.2 The juror's task

The notion of 'Reasonable Doubt' is famous for being the core concept of the highest standard of proof in Common Law systems: jurors in criminal trials are instructed to convict if and only if the defendant's guilt has been proven *beyond a reasonable doubt* (BARD). Although no clear and consensual definition of the BARD standard emerges from legal theory or practice⁸, one plausible option is to conceive of it as a probabilistic

⁶ Section 5 will develop on the variety of what accepting $non-H$ might imply.

⁷ As one anonymous referee suggested, one may also want to include some notion of partial acceptance (as distinct from acceptance of H -in-full); for example, if one's evidence is not great, one may take the lighter of two umbrellas. However, we suspect that this could be accounted for in terms of full acceptance by clearly specifying the decisional context: there may be one decision that concerns the heavier (and safer) umbrella, regarding which one does not accept H (because, given the respective costs of false negatives and false positives, one does not evaluate H as likely enough to carry the umbrella — just as for the decision whether to take one's plants out or not), and another one involving the lighter umbrella, regarding which one actually (fully) accepts H (concluding H is likely enough for one to decide to carry the lighter umbrella).

⁸ See *e.g.* Laudan 2008, 32-37, Roberts & Zuckermann 2010, 253-258.

threshold, by describing it as a function of the utilities associated with the various possible outcomes of the jury's decision, namely condemning / freeing an innocent / a guilty person.⁹ To be sure, establishing such a threshold is a far from trivial task: although Blackstone's (1765) formula according to which "*it is better that ten guilty persons escape than that one innocent suffer*" can be appealed to as an inspiration of such a consequentialist interpretation of the BARD standard, taking it (or any other similar consideration) literally would feel highly arbitrary. Besides, it is a debatable issue whether the BARD threshold should be thought of as fixed, or as a 'sliding scale'¹⁰, depending on the case (the seriousness of the crime, and relatedly the severity of the sentence).¹¹ Those are reasons to doubt whether it is possible to set a quantitative threshold (based on utilities) corresponding to BARD, and some have actually argued that we should drop such a view of the judicial standard of proof altogether, and rather interpret it as expressing some (high) level of subjective certainty.¹² As such, these are not objections to a DT view of 'reasonable doubt' outside the judicial domain: 'reasonable doubt' might ironically be trickier to define as a legal standard of proof than in other cases (although the very phrase 'reasonable doubt' stems from the legal context).

In this section, we nevertheless use the figure of the juror to explore some aspects of the DT account of reasonable doubt. We do not pretend to describe what actually happens in real courtrooms, no more than to explicate the BARD standard as jurors

⁹ Such a conception can be traced back at least to Justice Harlan's words in *In re Winship* (1970), 397 US 358, 363. See also Kaplan 1965, 1968, Allen 1977, Lillquist 2002, Laudan & Saunders 2009.

¹⁰ See Roberts & Zuckerman 2010, 253. See Picinali 2013 for a criticism of the DT approach to reasonable doubt based on the issue of the variability of the corresponding threshold.

¹¹ Another, related problem is that there seems to be something inconsistent in saying that the BARD standard is the highest standard of proof, used in criminal cases *because* stakes are particularly high in those cases, *and* to defining reasonable doubt as a threshold depending on the utilities. If one defines 'reasonable doubt' in BARD as set on the utilities, then BARD should apply to any case — 'balance of probabilities' corresponding to what counts as 'reasonable doubt' in the contexts of civil law. This might offer grounds for a criticism to DT accounts of the BARD standard of proof, and one might have to admit of two uses of 'reasonable doubt', one in terms of utilities, and one more specific as used in the BARD legal standard (thanks to Stephen John for helping us clarifying this).

¹² As a matter of fact, since the turn of the 21st century, the Crown Court has given up the very phrase "beyond a reasonable doubt" in its guidance to trial judges in directing the jury, in favour of a mention that "the prosecution must make the jury *sure* that [the defendant] is guilty" to prove guilt, and that "nothing less will do". (Judicial College, *The Crown Court Compendium – Part I: Jury and Trial Management and Summing Up* (February 2017), 5-8 (on line at www.judiciary.gov.uk/publications/crown-court-bench-book-directing-the-jury-2/). Note that such a restatement of the criminal standard of proof makes it considerably closer to the French rule of 'innermost conviction' (*intime conviction*).

should understand and implement it — or as it is intended by the law. Rather, we take (an abstract and idealised version of) the juror’s task as offering a simplified image of consequential decision-making under uncertainty: the hypothesis that is *in fine*¹³ at stake (guilt) is true or false absolutely (unlike, e.g., the hypothesis of the potential hazards of greenhouse gases¹⁴), evidence is strictly circumscribed, and only two decisions, irreversible and serious, are possible (acquit or convict)¹⁵. Further, an interesting feature of a trial verdict is that it is at the same time an assertion (supposed to express a proposition), and a consequential action (convicting someone both consists in declaring him guilty, and in sentencing him).¹⁶

Although the objections to the possibility of setting a threshold for BARD mentioned above are well-grounded, we want to highlight further issues with the threshold view of reasonable doubt, which arise even if we concede that some threshold can initially be set. This is relevant because in many everyday life situations, it does make sense to assume that there exists, at least in principle, such a threshold. Let us thus assume that a reasonable doubt threshold has been set for the juror’s decision task. This task can be schematically reconstructed as follows. Presented with a set of evidence of various types, the juror is supposed to evaluate the strength of this evidence, and accordingly update her estimate of the probability of guilt/innocence. Finally, she is supposed to assess the strength of such posterior probability against the supposedly fixed threshold, so as to bring a verdict. If guilt is confirmed beyond this threshold — if there is no reasonable doubt as to whether the defendant is guilty —, she must bring a *guilty* verdict. Otherwise — if there is a reasonable doubt —, she has no choice but to bring a

¹³ There are of course several hypotheses to examine on which the final hypothesis rests. See Schum 1994 for a thorough analysis of the “cascaded inference” at play in (legal) evidential reasoning.

¹⁴ Indeed, the question is not only whether it is true that greenhouse gases can be harmful, but how much. On the contrary, the accused is guilty or not of the charge — that being independent from the epistemic fact that the hypothesis of guilt is doomed, as any other empirical hypothesis, to be never 100% confirmed.

¹⁵ But see below footnote 45.

¹⁶ As suggested by one anonymous referee, this ‘verdictive’ aspect of verdicts might be a reason *not* to take the figure of the juror as paradigmatic. As our referee emphasizes, this contrasts with most other cases: for instance, accepting that climate change is anthropogenic might play a role in various practical decisions, but there is a gap between the acceptance and the subsequent action (of, say, imposing a limit on greenhouse gases emissions), which is often taken by another agent (here, governments). However, given our views on acceptance, and of assertion as a kind of acceptance (see 3.2), this makes the juror’s case even more interesting, by highlighting the essential link between acceptance, assertion, and action.

not guilty verdict.

1.3 The two cornerstones of the DT image

In the foregoing, we have assumed that the degree of confirmation for a hypothesis can be represented by means of a probability value. In the most standard DT framework, such probabilities are construed as agents' degrees of belief (hence as subjective, Bayesian probabilities). In the following, we will adopt the Bayesian framework: we will consider that the posterior probability of the hypothesis is the degree to which the agent believes this hypothesis after examining all available evidence.¹⁷ The DT image of reasonable doubt so sketched relies on two cornerstones:

[C1] For a given hypothesis H in a given decisional context, there exists a probabilistic confirmation threshold calculated according to the utilities, beyond which one should act upon H (*viz.* act as if H were true, as far as the decision at stake is concerned).

[C2] The only epistemic attitudes assumed by the DT account are graded beliefs (subjective probabilities). 'Acceptance' is not a distinct epistemic attitude: accepting H is acting based upon H , no more no less. In some situations, one may act upon a given hypothesis (accept it) even if one believes it to a low degree; in others, one may also decide not to act upon a strongly confirmed hypothesis.

Similarly, the DT framework does not say anything about doubt as a distinct mental state: 'reasonable doubt' appears as the complement to rational action (acceptance), and the DT framework does not say anything about the mental states one may entertain in the space of doubt. This does not amount to saying that accepting or doubting H do not come along with certain mental states, but the DT image does not imply this. The DT framework only posits graded beliefs and decisions based thereupon, the reasonableness of such decisions (be they acceptance or doubt) being assessed by reference to a consequence-dependent threshold.

Although the DT framework *prima facie* appears as an obvious candidate to account for important intuitions underlying the notion of 'reasonable doubt', this very notion does not seem to add in much to the DT analysis of rational action. In other words, the DT

¹⁷ In section 3, we clarify some assumptions and implications of such a view.

account of reasonable doubt is rather deflationary, ‘reasonable doubt’ being essentially defined in terms of graded beliefs and a probabilistic threshold for action. Is that all there is to reasonable doubt?

The remainder of this paper will be devoted to further exploring intuitions that underlie the notion of reasonable doubt as it can be appealed to in several situations, which will lead us to highlight some limits of the DT framework. Two main routes will be followed, which correspond to the two cornerstones abovementioned. After further analysis of the notion of doubt in section 2, sections 3 and 4 will question whether [C2] is tenable, by exploring whether all (seemingly) mental states at play in the space of doubt (beside graded beliefs) can really be described in terms of decisions whose reasonableness depends on utilities. Although the conclusion of this first journey will mostly be compatible with DT, section 5 will draw on this analysis to highlight some intrinsic limits of [C1].

2. More on doubt

In the previous section, we have seen that the DT framework, which seems *prima facie* well suited to capture essential aspects of ‘reasonable doubt’, is deflationary. Doubt is reasonable whenever the threshold for action hasn’t been reached; when it has been reached, doubt becomes unreasonable. However, in some situations, what we mean by ‘reasonable doubt’ seems not to be entirely captured by talking of the complementary to rational action. In fact, in situations such as the ones imagined in the trivial rain-and-umbrella example, asking whether it is reasonable to doubt would seem superfluously pompous: one would rather simply wonder whether to take one’s umbrella or not. In such cases, it is obvious that speaking of ‘reasonable doubt’ instead of rational action (based on reasonable belief) would add nothing to a DT account. But there are situations where the question of reasonable doubt seems to make much better sense. In such situations, is it always possible to reformulate the question of whether doubting a given hypothesis is reasonable in terms of whether one should act upon (accept) it?

In order to answer this question, we will start by deepening our analysis of the very notion of 'doubt'. Our goal, in this section, is not to specifically challenge the DT account, but rather to unravel aspects of the notion of doubt that underlies our study of 'reasonable doubt' as an enlightening category. As will be apparent, most of its aspects are compatible with the DT account. However, the clarification proposed in this section is a necessary preliminary to the remainder of the argument.

2.1. Doubt as decision-relative and dynamical

First of all, as finely accounted for in the DT framework, doubt is essentially relative to a certain decision. It seems to make little sense to ask whether one doubts a given hypothesis *tout court* (in absolute terms), independent from any particular decision to be made. Doubt does not merely correspond to lack of certainty.¹⁸ Nor does it systematically correspond to a low-degree belief. First, in some circumstances, a low confirmation degree can suffice to stop doubting. Second, we are far from certain of most imaginable hypotheses, but one would not say that one doubts them, simply because one does not care about them: if no decision is at stake, the question of doubt makes little sense.¹⁹ Be they likely or not, most hypotheses play no role in our reasoning and decisions: upon thinking of them, one may say that one believes some of them, disbelieves others, without the notion of doubt being relevant to account for those mental states. Is there somebody in the next room, from where you just happened to hear a sneeze? You hadn't wondered before, but yes, you believe so. However, if some important decision depended thereupon, you would probably start doubting it (*e.g.* imagining alternative scenarios leading to your having the illusion of hearing a sneeze from that room) and search for more grounds before accepting it. Conversely, if a difficult issue dissolves without one having to make one's mind regarding a hitherto

¹⁸ We do not claim that no occurrence of the notion of doubt in ordinary language corresponds to lack of certainty, but rather that the cases where it makes sense to ask whether doubt is reasonable are not of this kind. Our proposal is not intended as accounting for all the meanings that the word 'doubt' may have in different contexts; see Moon 2018 for such a proposal.

¹⁹ Such a stake-related aspect of doubt has been acknowledged by Thagard (2004), after Peirce (1877). However, it is not necessary, as Thagard does, to appeal to incoherence (according to him, doubt arises from the hypothesis' conflicting with one's background beliefs). Moreover, Thagard uses this idea to reject the probabilistic account of reasonable doubt; we however argue that they are compatible.

critical question, doubt may cease by itself, without making room for acceptance of any specific hypothesis regarding that question. In that sense, to stop doubting H does not always imply accepting H : one may stop doubting a hypothesis just because this does not matter anymore, without one's having to accept or reject it.

Consider a juror. If, at the end of the deliberation, she concludes that guilt has not been established beyond a reasonable doubt, she must bring a 'not guilty' verdict. Once this decision has been made, the question of doubt no longer arises: the issue is settled²⁰. Admittedly, the juror may still experience some doubt about the case, once out of the court. But such doubt is not related to the verdict and its stakes anymore. Moreover, assuming the defendant has been acquitted, the doubt she may feel should rather concern his innocence than his guilt: it would not make sense anymore for her to say "I doubt he is guilty". The situation has changed, so did the object of doubt, its stake, and the threshold for its reasonableness.²¹

As this example shows, the question of doubt specifically makes sense when the hypothesis at stake is per default accepted — be it the official one, the most widely held, or even the best supported and confirmed one.²² In other words, doubt is essentially dynamical, implying a somewhat effortful movement *against* the otherwise accepted hypothesis.²³ Imagine a scientist proposing a fanciful alternative to a theory in favour of

²⁰ This is the reason why it is not relevant, in the framework of the present analysis, to include in the definition of doubt acceptance of the complement to the hypothesis at stake (as does Salmon 1995 at the onset his paper): when one accepts *non-H*, it makes no big sense to say that one doubts H .

²¹ Certainly, as one anonymous referee emphasized, we can, and often do, have "idle doubts", namely doubts regarding a past decision, even though that decision is irreversible. That is however compatible with our view of doubt as decision-relative and dynamical: after the defendant was acquitted, the juror may doubt *whether she has made the right choice*. That hypothesis (let us call it H') is different from both the guilt (H), and the innocence ($-H$) hypothesis (which she may also doubt). The fact that she cannot come back on the decision on which it bears (which itself depended on H) does not imply that accepting or doubting H' has no relevance whatsoever for the ex-juror. Indeed, as the anonymous referee insightfully suggests, such alertness to the choices one should have made in the past may be a useful skill to cultivate if one wants to make better decisions in the future. One may account for the need to clarify whether one has made the right choice as an epistemic decision to be made, which might have practical consequences in the future.

²² True, presumption of innocence may suggest that the per default hypothesis should be innocence. But as soon as there is trial, it is clearly acknowledged that one should doubt innocence, and what becomes at stake is whether there is any doubt about guilt. This is one among other difficulties with how to account for presumption of innocence in a Bayesian framework (see Dahlman 2017, Fenton et al. 2017).

²³ In particular, raising the issue of whether it is reasonable to doubt makes particularly good sense when doubting would consist in actively questioning a hypothesis that the subject herself would be naturally inclined to accept.

which there is large consensus among scientists. It makes sense to ask whether it is reasonable to doubt the consensual theory and consider the alternative one, but it makes no real sense to ask whether one should doubt the alternative one (even though one may consider it with much scepticism).

2.2 A variety of possible routes in the space of doubt

Rather than a unique and stable epistemic attitude towards a given hypothesis, doubt thus appears to be the response to a question raised in a given decisional context. Such a response may imply a variety of distinct attitudes and actions. As we have seen, the juror on duty has no other choice, at the end of the deliberation, to either accept *H* or reject it. But, before the verdict, *in the space of doubt*, several routes might be followed, as remarkably illustrated by *Twelve Angry Men*. One may explore the prosecution's story so as to test its internal coherence and the plausibility of its various implications. One may also question the reliability of the sources of the evidence backing that story. Alternatively, one may try to build up exculpatory scenarios, etc. This is far from specific to the trial context. Quite similarly indeed, a scientist who doubts a hypothesis that is consensual, or even unanimously endorsed in the scientific community, has a great variety of possible search strategies. Among other options, she may explore this hypothesis, with the hope of showing its flaws by drawing obviously unacceptable conclusions from it, or else formulate an alternative, incompatible hypothesis, and explore its consequences so as to show its fruitfulness. Depending on the time and cognitive resources she is able to invest, she may also do both in parallel.

Although the DT framework does not say much about the variety of possible reasoning routes and actions in the space of doubt, the conception of doubt we have just sketched is so far compatible with the DT framework.²⁴ In the next two sections, we will however explore two ways we might actually be tipped out of such framework.

Consider the juror out of the court again. First, as we have said, even though she has brought a 'not guilty' verdict and the defendant has been acquitted, she may still be

²⁴ As suggested by an anonymous reviewer, the general picture of inquiry sketched here could be nicely complemented analyses of the relationships between Inference to the Best Explanation and Bayesian accounts.

experiencing some doubt regarding his innocence. What would such doubt imply, and to what decision would it be related? What makes it (un)reasonable, if anything? We will address those questions in section 4, exploring whether there could in some situations be reasons (not) to doubt that do not relate to any decisional context, and do not reduce to a cost-benefit calculation, and shedding light on the reasonableness of the various possible routes in the space of doubt.

Before then, let us consider this second possibility: rather than doubting his innocence, our juror might actually *believe* that the defendant is guilty (even though she didn't accept this in the trial context). In the next section, we will temporarily leave the notion of doubt aside, and try to clarify what that would mean, and whether it is possible to account for such a situation in DT terms.

3. 'Accepting' *non-H* while 'believing' *H* — a detour

Consider our juror: once out of the court, where she brought a verdict of acquittal, she confesses to a friend — or even to herself — that she actually *believes* that the defendant had committed the crime he was accused of. She does not merely say that she holds a certain degree of belief in the hypothesis of guilt (and a complementary degree of belief in the hypothesis of innocence), but rather that she somehow *has* the belief that he is guilty, and *does not have* the belief that he is innocent. She may certainly be more or less confident, as well as more or less prone to change her mind (her belief does not necessarily amount to certainty), but whatever its strength or entrenchment, her belief does not seem to correspond to the same kind of mental state than graded beliefs: even though it can be qualified, the expression of such belief takes the form of a 'yes-no' answer to a categorical question, rather than a graded answer to a 'how much' question. In other words, one either believes *H* or not.²⁵

²⁵ Not believing *H* may imply having no idea, but that doesn't correspond to a degree of belief either.

How can such notion of belief — and the epistemic situation of the juror who accepts *non-H* while believing *H* — be accounted for in the DT framework?²⁶ One answer would be to consider such beliefs as non-(entirely)-rational, passive psychological states, which do not fall within the intended scope of DT.²⁷ Another one would consist in acknowledging that such beliefs could be rational and evidence-based, thus requiring some amendment to the DT framework (at least as we presented it in section 1), by positing other mental states beside graded beliefs. After briefly considering the first option, we will explore the second, which will lead us to finally propose that ‘full’ beliefs may be identified with some specific kind of acceptance — thus suggesting that they might also be context-sensitive.

3.1 *Belief and acceptance as two different regimes?*

A first way to describe our juror’s situation is to treat belief as a passive psychological state, whose formation would be beyond rational control, even though it might partially be influenced by evidence. Acceptance, on the other hand, would result from one’s attempt to reason following some rational procedures — which are sometimes provided as more or less explicit rules (*e.g.*, legal rules of evidence, standards of scientific methods and proof, etc.).²⁸ Such a distinction would roughly correspond to Cohen’s (1995) own distinction between belief as a passive and non-rational mental state on the one hand, and acceptance as an active, voluntary, and reasoned endorsement of a claim. And in fact, Cohen’s (2002) claim that jurors are asked what they accept, rather than what they believe, is based on this distinction. This view would also match Lackey’s (2007) example²⁹ of a ‘racist juror’ who cannot shake his belief that a black man accused of having raped a white woman has actually committed the crime, while acknowledging

²⁶ Note that our view of acceptance is compatible both with cases of belief without acceptance, and with cases of acceptance without belief. In fact, as an anonymous referee emphasized, the notion of “acceptance” in the philosophy of science literature is most typically used as a way of understanding cognitive attitudes which, in some sense, fall short of “full belief” (see Jeffrey 1956, Levi 1960).

²⁷ They would not be related in any way to rational action.

²⁸ Thanks to Stephen John for helping us clarify this.

²⁹ The goal of Lackey in that paper is to show that norms of assertion are different from norms of belief (that it is neither necessary, nor sufficient to believe *x* in order to be entitled to assert *x*).

that the evidence does not support such a claim. Being aware that his non-discardable belief is prejudice-based, the juror tells others that the defendant is innocent (and probably, a fortiori, brings a verdict of acquittal, though Lackey does not make this explicit).

There are several problems with such a view, among others the difficulty of understanding the relation between such a resilient, non-rational state, and the epistemic states underlying acceptance (it is indeed hard to account for acceptance so construed as a first order mental state). In any case, admitting such a view implies a renouncement of the Bayesian DT framework altogether. Discussing these issues would require a long development, which would be beyond the scope of this paper. We will rather try to account for this by sticking to the DT framework.³⁰ Moreover, even if such passive and irrational psychological states not accountable in the DT framework existed, what we aim at capturing here is different. What we want to account for is the situation where a juror rejects H (hence accepts *non-H*³¹) while *rationally* believing H . Because the standard for finding someone guilty has to be particularly high, such a situation does not seem unrealistic. As should be clear from the preceding analysis, accepting *non-H* is nothing more than making a certain decision (taking a certain action) — in this case, bringing a verdict of acquittal.³² This does not imply any other particular mental state

³⁰ Supposing that the case imagined by Lackey correspond to a real situation, a Bayesian account would consist in considering the juror's belief (which he refuses to base his verdict on) as actually updated in front of the evidence presented in court, but starting from a pathologically high prior — or that his evidence evaluation is biased by racist prejudices and that he gives more weight to incriminating evidence. Being aware of that, the juror would correct this bias by the establishment of a higher threshold for reasonable doubt. In fact, the DT framework is able to account for prejudice-based priors and wrong evaluation of evidence. From that perspective, having explicit rules of evidence, and basing verdicts on collective decision, as well as the practice of peer-reviewing and standards of scientific proof, would in fact be a way to compensate for the irreducible subjectivity of belief.

³¹ Indeed, at the end of the day, jurors only have two options. There is certainly time for deliberation, but, when the verdict is brought, rejecting H implies accepting *non-H*.

³² However, acceptance as we think of it has to be based on evidence. Cohen, who emphasises the difference between belief and acceptance, is surprisingly unclear about that: in order to emphasise the voluntary character of acceptance as an action, as opposed to the passivity of belief, he claims that a juror bringing a verdict of acquittal because he sees the penalty as “being too severe”, or the alleged crime as “an honourable deed or a public service” (Cohen, 2002, 298-299), “would then be accepting, on ethical or pragmatic grounds rather than on cognitive ones, that the accused is not guilty, while believing perhaps quite firmly in his or her guilt. So the verdict would declare what the jury accepts; not what it believes.” However, if the juror does so independent from any consideration of the evidence, then he does not *accept* the hypothesis of innocence in our sense — the question of guilt and its probability is simply not relevant to his verdict. In divorcing acceptance from belief, Cohen misses the evidential grounding of the latter.

towards the defendant's innocence than a certain $p(H)$, which may or may not give rise to acceptance, depending on the decision at stake.³³ Although this might feel uncomfortable, it does not sound irrational or contradictory to estimate that the standard for conviction hasn't been reached (hence bring a verdict of acquittal), while nevertheless silently believing that the defendant has actually committed the crime he is accused of — and believing it with some reasons.

3.2 Categorical belief and sincere assertion — the context-dependence of belief

The picture of mental states proposed by Maher (1993) seems likely to account for the situation described above.³⁴ According to him, full belief (which he calls “acceptance”) is different from both graded beliefs and rational action. Unlike the latter, it is not a response given in a certain decisional context; full belief is a mental state, independent from any specific decision. However, it can be rational and evidence-based. Maher defines such full belief (‘acceptance’ in his terms) as the object of “sincere and intentional assertion”. Although he does not explicitly include assertions that are

³³ The verdict's taking a linguistic form (typically the form of an assertion about guilt) may be misleading, but the question asked to jurors is whether guilt has been proven beyond a reasonable doubt — not whether they ‘believe’ that the defendant is guilty (as Cohen rightly acknowledges, although we disagree with his definitions of belief and acceptance). Interestingly, as a way to prevent from confusing rejection of guilt and *belief* in innocence, juries in Scotland have a choice between three verdicts: ‘guilty’, ‘not guilty’, and ‘not proven’, the last two equally imply acquittal. Although this can be thought of as a protection of the defendant, this was heavily criticized (Walter Scott called it the *bastard verdict*), as stigmatizing defendants for which the verdict was ‘not proven’, since they are not freed from any suspicion.

³⁴ There exist several attempts at conciliating a graded view of beliefs (especially in the Bayesian perspective), as being the most likely account of rational belief change, with the notion of full belief, as both highly intuitive and corresponding to the folk psychology view of mental states and the representationalist view of mental life, as well as largely assumed by most mainstream epistemology. Those attempts have given rise to a debate in epistemology and the philosophy of mind about the relation between full beliefs (also called ‘categorical’ or ‘outright’ beliefs) on the one hand, and graded beliefs, on the other (see Ebert & Smith, eds. 2012, as well as the online bibliography available at <https://natureofrepresentation.wordpress.com/2013/09/27/bibliography-on-belief-and-degrees-of-belief/>). For a more formal approach of the question by a Bayesian who nevertheless keeps room for full belief, see Leitgeb 2013. This debate is sometimes practically undistinguishable from the discussion on the relations between belief (categorical or probabilistic) and acceptance (see Engel, ed. 2000) — those notions being diversely construed by different authors — and on the norms of assertion (for a recent contribution, see Hawthorne, Rothschild & Spectre 2016). We will not present those debates, and the various arguments pro and contra the reduction of full beliefs to graded beliefs (or vice-versa). We nevertheless take for granted that full belief cannot be equated in any simple and general way to a specific degree of belief, whether very high, or simply above 0.5.

constrained by explicit rules such as verdicts, and he may not agree to count scientific claims as being also constrained by such rules, hence as being essentially different from “sincere” assertions³⁵, his view seems likely to match ours: most assertions would count as acceptances whose reasonableness is consequence-dependent, but there would remain a subset of assertions that would be perfectly sincere, and not context-sensitive — hence essentially different from acceptance of any sort. Let us now question such a view.

What is sincere and intentional assertion? Imagine a friend of yours asks you whether the chocolate cake you’ve just baked contains peanuts. You know you haven’t put any peanuts in it. You would happily answer that it is peanut-free. This is a perfectly sincere and intentional assertion, which expresses your belief that the cake is peanut-free. However, if you know your friend has peanut allergy, you may answer you’re not sure about it — after all, you haven’t really paid attention to the exact composition of the chocolate you’ve been using.³⁶ This is not a less sincere and intentional assertion than the previous one. Even such a non-procedural verbal exchange seems to obey some consequence-dependent standards. As a consequence, from a highly constrained context such as the trial one, to the most familiar verbal exchange, there seems to be a continuum, where confirmation thresholds for assertion vary depending on the epistemic relation between interlocutors (doctor-patient, expert-judge, expert-policy-maker, epistemic pairs, etc.), and on what is at stake.³⁷ From this point of view, assertion is, just as other actions, context-dependent.

Is there something more to “sincere and intentional assertion” than such context-dependent speech acts? Defining ‘sincere and intentional assertion’ as the response to a decontextualized question, with no further consequence than belief expression, would be obviously circular. Let us try, however, to imagine an opinion poll questionnaire that would be taken in a perfectly neutral context — if that is possible at all. Even if we

³⁵ On sincerity, and how it is not the right value for understanding scientific claims and many speech acts, see (John 2018). Although our view is compatible with this claim, it makes finally hard to grasp what a ‘sincere’ speech act would amount to, as we will shortly see.

³⁶ Maybe you would answer “I believe it is peanut-free”. But this use of ‘believe’ is not so much a description of your mental state, as a way to qualify the strength of your assertion. See Simons (2007) and Macready (2014) for a pragmatic analysis.

³⁷ Moreover, it would be a naïve idealization not to acknowledge that one’s attribution of mental states to others impacts one’s conveying one’s own mental states (even sincere).

assume that subjects make a sincere introspective effort, nothing allows us to assume that they have an easy and transparent access to their mental states³⁸, even more so if they have to express them linguistically. The way questions are posed, and the very format of the expected answer, are probably far from neutral.

Back to the chocolate cake and allergy example: isn't it possible to say that you believe your cake is peanut free, but that you would rather not assert it? In that sense, there would be something like context-independent beliefs; sincere assertions expressing them would rather be an idealization, or at least would generally come together with explicit qualification according to the context ("I don't believe there is any peanut, but I cannot guarantee it").

Maybe there are such (context-unaltered, but most often silent, and possibly inaccessible) beliefs: as a widely shared picture of mental states suggests, there would, for each individual, be something like elementary "bricks" entering into her "belief box". But how we access them is a far from trivial issue. Finally, we do not see how one could distinguish them from certain responses, given in certain contexts. Full belief, considered independently of any action eliciting it, is a hard to grasp entity, much harder to grasp than the usual notion of belief, and with it the epistemological analyses that take it as the fundamental material of our mental life, suggest. The closest thing we have access to, has a context-dependent threshold, and therefore corresponds to a sort of acceptance, based on a specifically graded belief.³⁹ In that sense, the juror's belief in the defendant's guilt, since it is expressed (even though in a private sense, as the conclusion of an introspective process), corresponds to an act of acceptance, with a different threshold from the verdict's one.⁴⁰

³⁸ In fact, many philosophers and psychologists claim that we do not access to our beliefs through introspection. See Mandelbaum 2014, p. 75-76 for references.

³⁹ The views defended here are close to so-called 'pragmatic encroachment' positions in epistemology, which argue that knowledge is 'interest-relative', and that knowledge ascriptions should be context-sensitive (see Weatherson 2005, Ross & Schroeder 2012, Wedgwood 2012).

⁴⁰ It is a hard question what actually determines the threshold for full beliefs so construed, and we can only draw conjectures about empirical issues. One can for instance suppose that the juror's "full belief" is not independent from the fact that she otherwise has a verdict to bring. Maybe her threshold for belief would be high, so as to avoid cognitive dissonance; conversely, she might react to the uncomfortable feeling of bringing a verdict of acquittal while being far from convinced of innocence by acknowledging (at least for herself) that she believes the defendant is guilty... In all those (putative) cases, the threshold for reasonable doubt as far as the verdict is concerned would not be free from consequences on the threshold for what the juror identifies as her "belief". More generally, it seems hard to ask oneself what

Although we don't pretend to have demonstrated that there is nothing like full belief, which would not reduce to any kind of contextually defined degree of belief — that would be beyond the scope of this paper —, we haven't found a compelling reason to amend the DT framework to accommodate a categorical (as opposed to graded) mental state that would not be context-dependent. In what follows, we will therefore consider that so-called 'full beliefs' are nothing but a kind of acceptance. This detour therefore serves the purpose of clarifying how much we are ready to buy from the DT framework — namely the essential context-dependence of anything that may appear as a stable, 'full' (in the sense of non-graded) mental state — so as to make the limits of our final criticism precise. It will also offer a basis for the argument in the next section.

4. Back to doubt: aspects of its (un)reasonableness

Let us now come back to the first problem raised at the end of section 2: wouldn't it sound sensible to say that, once out of the court, the juror does not have any further decision to make for which the defendant's guilt may be relevant, and that the doubt she may experience is not relative to any specific decisional stake? If that was the case, it would be difficult to account for the reasonableness of doubt in terms of utilities.

First of all, it is worth noting that there are decisions for which whether the defendant has committed the crime matters, even out of the court's context. Our juror might for instance doubt guilt with regard to the trial context (estimating that the threshold for

one believes, without considering any of the decisions that are likely to result from the answer, or without imagining a possible interlocutor. Not only is it difficult to imagine an entirely non-contextual assertion, but also can we hypothesise that, for what might be the closest to such an ideally sincere assertion, the threshold for belief might vary along various parameters. Among other things, what (one believes) others believe about a given issue might make one's own answer vary. The threshold might also vary according to the threshold for a decision that is typically associated to a question (which may vary between subjects). For instance, is it really possible to consider the question of what one believes regarding the potential harmfulness of vaccines without the debate on obligation to vaccinate, and the various opinions expressed around this entering into one's assessment of one's own belief? Is it possible to entirely dissociate the issue of the guilt of a defendant from the typically associated decision that is verdict — even though one does not have to actually bring a verdict? Those are psychological issues, whose answer can only be empirical, but raising them here enables us to highlight that full belief, as distinct from acceptance, is a very hard to grasp entity. The closest thing we have seems to vary depending on the context, just as any other act of acceptance.

conviction hasn't been reached), but still doubt his innocence with regard to continuing to live in a town where this potentially dangerous man lives. Her doubt might therefore prompt her to move home. To take a less directly practically-orientated (and dramatic) decision, we have seen in the previous section that most acts of assertion may count as context-dependent acceptances: our juror may doubt guilt (and accept innocence) in court, but she may doubt innocence with regard to what to tell others (family, friends, or maybe journalists — supposing this was allowed to jurors). But if there really is no decision of any sort at stake, if our juror is simply unsure, with no further consequence, it would not make sense to say that she 'doubts': she rather is in a passive state of belief suspension, where doubt no longer arises, because no specific question is asked.

If, however, she is not sure but is not content with such a state, and wants to know more about the case, just for the sake of it, then she may arguably say that she "still doubts". But how could the reasonableness of such doubt be assessed? In fact, since no empirical hypothesis is ever 100% sure, how could such a 'purely theoretical' doubt be considered unreasonable at all? If she considers that she needs to know more to be able to peacefully stop doubting, how could one object to her that this is unreasonable? Maybe after all are we not even entitled to qualify of 'unreasonable' someone who would doubt that the Earth is round from a purely theoretical point of view (if no practical consequence follows)? Similarly, if no practically detrimental decision follows, doubts about the rotundity of the Earth, or about Al-Qaeda's being responsible for the 9/11 plot should not be called unreasonable?

First of all, one must note that the frontier between purely theoretical and practically orientated inquiry (or doubt) is not clear cut, and that there rather seems to be a continuum from the most practical to the most theoretical. Although moving house is clearly practical, and continuing the inquiry for the sake of it seems purely theoretical, what about, for instance, judging others (their reliability, expertise, the quality of their information sources, etc.) based on their opinion about a certain question? Doubting a claim implies doubting the sources thereof, and the credibility of people who endorse it. Such a 'chain reaction' might be practically consequential: if you start doubting a given claim that is unanimously endorsed by the scientific community, you may finally end up doubting the reliability of vaccination, for instance. And in fact, as highlighted in

the introduction, many conspiracy theorists or ‘merchants of doubt’ primarily appeal to ‘critical thinking’ and ‘open-mindedness’, which are supposedly harmless (at worse, and, most often, beneficial).

However, even though one abstracts from any practical consequences that may follow from doubting a claim, it seems that ‘purely theoretical’ doubt could possibly be considered unreasonable in some other sense(s), which have to do with allocation of cognitive resources. Consider our juror, still doubting the defendant’s innocence. What would such doubt imply? From just keeping her mind open to the hypothesis of guilt, thinking of it and reconsidering evidence to getting passionate about it, leading her own inquiry by systematically exploring several scenarios of what happened, and even trying to gather more evidence, she might invest more or less time and energy in this issue.⁴¹ Independent of the practical consequences of such an investment (if really passionate with the issue, she might neglect her usual business, family, or even do so at the expense of her rest and health, etc.), whether the purely *cognitive* costs are worth spending depends on several parameters, including the importance she gives to this issue relatively to others, at the expense of which she allocates her cognitive resources to it, but also the chances she has to actually find new information that may make a significant difference.⁴² The question of the best allocation of cognitive resources therefore also includes an estimation of the subject’s own competences in comparison with others’: on some subjects, one should rather defer to others⁴³ (more experts on the matter), both because they are more likely to find and adequately process relevant information, but also because one’s own cognitive energy would be better allocated elsewhere.⁴⁴

⁴¹ As one anonymous reviewer notes, there should be a way to develop our account of doubt as essentially decision-relative in a way that makes room for a general disposition to doubt — by contrast with active doubting — as valuable, insofar as it warrants openness to potential alternative hypotheses.

⁴² This argument is different from Huneman’s (2015), who suggests that adopting some conspiracy theories might result very costly from an epistemic point of view, because of the radical revision of one’s web of belief it implies. This approach rather applies to the *adoption* of fanciful theories than for doubt. Indeed, doubting a theory does not imply accepting another one. Hence, doubting well entrenched theories does not necessarily lead to incoherence. At any rate, our argument is not incompatible with Huneman’s.

⁴³ See Hardwig 1985.

⁴⁴ One could even add some global coherence constraints on thresholds, whereby one should doubt any claim that is less confirmed than a given claim I doubt. However, such constraints seem to conflict with others, having to do with social norms. For instance, imagine a person who has no a priori reason to

Those considerations drawn from consideration of our fictional juror have obvious consequences for scientific questions of societal importance. Considering them shows that other grounds for judgements of (un)reasonableness might be found in some value judgements. Optimization of collective inquiry by means of division of cognitive labor could conflict with thirst for individual knowledge. On a different level, a certain conception of epistemic responsibility could prompt someone to escape from a road often taken by conspiracy theorists and “merchants of doubt”, likely to pulling one away from one’s own landmarks and finally to leading one to give up the very epistemic values that initially made one doubt. Conversely, a certain conception of intellectual courage and honesty could prompt one to take it.⁴⁵ Clearly, depending on all those parameters, judgements of the (un)reasonableness of doubt must vary from individual to individual: not only do different people not have the same evidential basis and competences, but they do not necessarily share the same values.

The issue of the allocation of cognitive resources appears in a particularly salient way in the case of ‘pure’ theoretical doubts such as our ex-juror’s. Indeed, even if one assumes that there might be some kind of confirmation threshold beyond which she would be satisfied and stop doubting, one can also think of theoretical inquiry as potentially never ending: unlike other cases where decision has to be made at some point for some practical reasons, there would be no reason to settle the issue beside problems of cognitive resources limitation. However, all the aspects of the (un)reasonableness of doubt highlighted in this section are not essentially different depending on whether there is something practical at stake or not. Indeed, even in inquiries that are, *in fine*, highly consequential, search strategies *in the space of doubt* are certainly constrained

believe that she has a certain disease that is relatively rare (let’s say a 1 on 3 000 incidence in global population). If that person decides to take the diagnostic test that no one takes unless one has some reason to suspect that one is ill, one would think this person is unreasonable (she unreasonably doubts that she is healthy to this respect and search more evidence). Now imagine another person who is considered at risk for an extremely rare disease: the incidence in general population is 1 on 1 000 000, but that person has 1 on 3 000 chance to suffer from this disease. She will be considered totally reasonable to take some test to know whether she has this disease (even though we assume same costs for test, and same gravity of disease, and equal opportunity to cure it).

⁴⁵ Alice Dreger’s inquiries, as she tells them in *Galileo’s middle finger* (2015) embody this kind of intellectual (and more general) courage. This leads her to deep reflections on whether any hot question is worth debating, raising both the question of whether one should address any question whatever one may find (even if that might somehow be detrimental for social justice) but also because one may initially think that there is no chance to find anything on such a route.

by time and resources (of different kinds) limitation. But this does not seem essentially different from purely theoretical inquiry as paradigmatically exemplified by the character of Sherlock Holmes, even a highly consequential and practically orientated quest such as criminal inquiry might resemble in many ways a purely theoretical one, where one may try several strategies, among others exploring some initially implausible route.

Where are we now? Despite their complexity and multidimensionality, the various parameters entering the evaluation of the (un)reasonableness of doubt that we have presented in this section seem all to be capturable in terms of allocation of resources. Allocation of resources itself should be describable in terms of costs and benefits. It must already be obvious that factoring in those costs and benefits in estimating whether the threshold for acceptance (hence the threshold beyond which one should just stop doubting) has been met is far from trivial. Given the variety of possible routes in the space of doubt, and the multidimensionality of the parameters entering the costs and benefits calculation, one has reasons to suspect the threshold view is too coarse a simplification to account for reasonable doubt. Let us now come back to the threshold view in the light of the above analysis.

5. Reasonable doubt as a probabilistic threshold: some limits

Let us briefly take stock of where we are. In section 1, we have highlighted two aspects of the DT account of reasonable doubt, namely the idea of a probabilistic threshold (C1) and the assumption of graded beliefs as only epistemic attitudes (C2). Deepening our intuitions regarding reasonable doubt, we haven't truly challenged C2, but we have highlighted the complexity and multidimensionality of the (un)reasonableness of doubt, of which the DT framework as such does not say anything. Although all these aspects can in the end be described in terms of utilities, their articulation with the threshold view needs to be explored. To do so, we will highlight two problems that make the idea of a probabilistic threshold too narrow to account for the notion of reasonable doubt.

Those two problems, although distinct, both have to do with how we deal with our ignorance, and especially our ignorance of the extent of what we ignore.

As mentioned in section 1, not accepting H might lead, depending on the situation, to accepting $non-H$ (in which case doubt no longer arises), or instead to seeking more information before accepting or rejecting H . However, what $non-H$ is needs to be clarified, and a distinction needs to be drawn. In some situations, $non-H$ needs to be made specific (a specific alternative hypothesis has to be proposed, and accepted in order to reject H). In others, one needs not necessarily make $non-H$ specific to reject H . We will alternatively consider problems arising with the two types of cases.

5.1 'Unpacking' $non-H$

As mentioned earlier, jurors have no choice, if they do not accept H , but to accept $non-H$. However, the DT framework does not say much about what $non-H$ may be, beyond the negation of H . Given what we have seen in the previous sections, as to the various possible routes in the space of doubt, we however suspect that an account of reasonable doubt in terms of a probabilistic threshold is not enough if it does not take into account the importance of $non-H$ in the evaluation of whether the confirmation threshold for H has been met.⁴⁶ In other words, there might be reasons to doubt H that are not well explicated by merely describing them as reasons not to accept H (as absence of sufficient reasons to accept it).

Consider a juror again. In deciding what verdict to bring, she is instructed to only use the evidence that is presented in court (hence the question of whether to gather more evidence does not arise). The evidence she is supposed to rely on to calculate the posterior probability of guilt may vary in quantity as well as in quality. As a result, this probability is more or less robust and resilient to potentially additional evidence. Suppose that, on the basis of the evidence she was presented with, a juror evaluates the probability of guilt to a very high degree — at any rate sufficient to meet the reasonable

⁴⁶ This issue is related to another one, which we will not address here, about whether statistical evidence is enough to convict — such as the so-called “gate-crasher paradox” (Cohen 1981).

doubt threshold. She can however suspect that the inquiry wasn't carried out properly. This does not merely mean that some pieces of evidence might have been fabricated, or that some witnesses may be lying. Part of the jurors' task is actually to evaluate witnesses' reliability and evidence overall credibility. Let us thus ignore that aspect, and suppose that all the evidence presented in court is credible and relevant (or that relative credibility has already been factored in). It remains possible that, among the elements that were *not* presented, some would have been relevant — at least from our juror's point of view⁴⁷. More generally, inquiry is doomed to be partial. In some cases, a juror may suspect that other items could have been collected, which might counterbalance the evidence actually presented in court. In this case, she may consider that the confirmation degree of the hypothesis based on the available evidence, however high, might get significantly lower were further pieces brought. In other terms, she may be more or less confident in her own probabilistic judgement. It seems clear that this must be taken into account when assessing whether the reasonable doubt threshold has been met. Hence, whether it is reasonable to doubt a hypothesis seems to depend on second order, rather than on first order factors only.⁴⁸

Although the DT framework is certainly able to integrate such second order factors, this may not be as trivial a task as it may seem. To understand why, let us draw a distinction between two types of missing evidence. On the one hand, there is evidence which one knows is missing; for instance, if the inquiry shows that a car is implied in a crime because there are tire tracks indicating it, but has no clue as to what the brand of the car is. That is information that would be relevant for the inquiry (and *in fine* for evaluating the guilt of a particular person), but which could not be gathered for some reason. Clearly, a good Bayesian reasoner should integrate in her final degree of belief her own assessment of the quality of the evidential basis (hence her assessment of how much relevant evidence is missing). On the other hand, there is all the evidence that one does not have because it would constitute evidence for hypotheses inquirers haven't explored, considered, or even formulated — but which might be relevant as alternative

⁴⁷ See Schum 1994 on the irreducible subjectivity of the evaluation of the relevance of evidence.

⁴⁸ As emphasized by Stephen John in personal communication, this is well exemplified by debates over climate modelling: one can have some degree of belief in some claim generated by a model, but also have reasons to worry the model itself is partial.

hypotheses to H — in other words, as specific *non-Hs*. As soon as there is a particular suspect, criminal inquirers generally focus on the collection of evidence that is relevant for the evaluation of the suspect's guilt (be it incriminating or exculpatory evidence); even though they may leave some other options open, they cannot gather all the evidence that might possibly be relevant for entirely unconsidered scenarios. Once the case is brought to trial, the guilt hypothesis — let us call it H — is made specific; it gets fleshed out as a more or less detailed scenario (minimally, in the case of a murder, “the defendant murdered the victim”). On the other hand, the innocence hypothesis — *non-H* — often remains unspecified. Admittedly, the defence can propose an exculpatory scenario — hence a specific *non-H₁* (e.g. by advocating the thesis of an accident), but that does not exhaust all scenarios satisfying *non-H*. What would happen to one's degree of belief in H if some new evidence relevant for a particular *non-H* were brought is a difficult question, but one obviously important for any account of reasonable doubt as a probabilistic threshold.⁴⁹ How can one account for whether the reasonable threshold has been met based on reasons to accept H , if one does not consider potential reasons to accept specific *non-Hs* — hence reasons to doubt H that do not reduce to reasons to reject it? Clearly, all the potential routes that a free inquirer as the tireless former juror imagined in the previous section could take (which the juror on duty cannot), and what they might lead to, should play a role in estimating whether the reasonable doubt threshold has been met.

3.2 Infinite regress of ignorance and doubt

The juror's situation makes the problem exposed in the previous subsection particularly salient and clear. Indeed, jurors are not in charge of the collection of evidence; moreover, among the evidence gathered in the course of the inquiry, only the pieces that are judged admissible are presented to the court, and jurors are explicitly instructed not to take any other piece of information they may possess into account. However long the deliberation may be, at the end of the day, the juror must conclude on the exclusive basis of the elements presented in court; she cannot pursue the inquiry by herself — or

⁴⁹ This is object of Hahn and Hartmann, in progress.

even ask for more elements to be brought to the court. Finally, even though one does need to come up with a specific *non-H* to reject *H*, there *must* be a specific *non-H* explaining what actually happened if *H* is not true.

However, in many other situations, settling the issue does not necessarily amount to finally accepting a specific hypothesis (either *H* or a specific *non-H*). This may give rise to another kind of problem, which also makes the idea of a probabilistic threshold for reasonable doubt too simplistic.

Let us call *H* the hypothesis that a patient suffers from disease *D*, which only medication *M* may cure. The decision at stake is whether to take *M* or not: in this situation, accepting *H* implies taking *M*. Here, not accepting *H* at time *t* does not necessarily imply definitively accepting *non-H*; rather, one can consider that one does not know enough yet to accept *H*, and search for more information before settling the issue. Moreover, unlike in the juror's situation, the question of specific *non-Hs* is not as pressing: the urgent question is whether one should take some medication to cure *D* (hence whether one actually suffers from *D* or not), but not necessarily what the right explanation is. There might indeed be nothing special to explain: one may be in a risk group for some disease, without one's having any particular symptom that needs to be explained otherwise if it turns out that one does not suffer from that disease. In other words, if *H* is false, that might not necessarily imply that *non-H* needs to be specified; it might be sufficient to reject *H*.⁵⁰

On current knowledge, the confirmation threshold of *H* (for accepting it) depends on several parameters, among which the seriousness of the consequences of *D* for the patient, the side-effects of *M*, and its chances of success, but also the risks and potential benefits of delaying the decision. What could be the potential benefits of delaying the decision — hence doubting *H* without accepting *non-H* and settling the issue? First of all, one may of course gather more information relevant to the degree of confirmation of *H* by undergoing further analyses. Second, one's doctor may also make some progress in her knowledge of *D* (if *D* is a rare disease on which little research has been done, and

⁵⁰ This is far from a rare situation. Most application of the precaution principle consists in acting as if *H* was the case, without necessarily searching an explanation for a specific phenomenon.

is still rather confidential, it is possible that a few additional studies, or a more thorough inquiry, suffice to elaborate other treatments, or to uncover serious side-effects of M). As mentioned in the previous section, one crucial question is whether we have any chance to change the confirmation degree in a significant way. But, in deciding whether to doubt H or not must enter considerations about one's chances to find new information that might modify not only the confirmation degree of H , but also the threshold for accepting H itself; indeed, information about the evidence regarding H is not always independent from evidence about the potential consequences of accepting H . All these potentialities of changing one's current state of knowledge should be probabilistically assessed, and they are so on the basis of one's current state of knowledge. Hence infinite regression threatens. In other words, delaying opens up a space of potentiality that make the problem of potentially infinite complexity.

Hence, the threshold view is simplistic for another reason: the question is not only how much H is confirmed and how robust this degree is, but also how robust the setting of the threshold is. Against the risk of delaying, one must therefore evaluate one's own ignorance, not only of what we know, but also of what we ignore.

The two problems addressed in this last section show that the issue of reasonable doubt has to do with how we deal with our ignorance, and how we factor it in in our calculation. Without going any further, let us emphasise how focusing on doubt leads to raise issues that the usual notions of belief or acceptance tend to obscure.

Conclusion

We hope to have shown that a doubt-centred approach reveals the richness of the epistemic attitudes and actions (including purely mental actions such as reasoning processes) at stake in any decision-making, which are often obscured by belief-centred approaches. Rather than what one *believes*, the question of reasonable doubt concerns what one *ignores*, and how one manages one's own ignorance. The latter being unbounded, and undertermined, the problem is multidimensional, and of a virtually infinite complexity. The way each one decides to allocate one's resources (both

cognitive and other) depends on a complex web of epistemic attitudes and values. The intuition underlying the DT picture, according to which the limit of reasonable doubt depends, *in fine*, on each one's preferences, may be right. But the ideas of a probabilistic threshold set by the context and of graded beliefs appeal for some enrichment by a deepened analysis of possible mental states and actions, which we hope to have sketched here. In doing so, we hope to have proposed a plausible integrated account of what makes it possible to state that the stubborn juror's scepticism is reasonable, while the 9/11 official's theory denier's isn't, by at the same time questioning the idea of a clear-cut distinction between theoretical, and practical aspects of doubt.

Finally, let us mention that no analysis of reasonable doubt could be complete without integrating a dimension essential to any epistemic decision, which we had to ignore here, namely its collective aspects. Like a juror in a jury, what (one believes) others believe plays a determining role in one's own epistemic decisions — which increases even more the complexity of the problem.

References

- ALLEN, R. 1977. The Restoration of *In Re Winship*, 76 Mich. L. Rev. 30, at 47
- BLACKSTONE, W. 1765. *Commentaries on the Laws of England*
- BONJOUR, L. & SOSA, E. 2003. *Epistemic Justification. Internalism vs. Externalism, Foundations vs. Virtues*, Malden, MA: Blackwell
- BRATMAN, M. E. 1992. Practical Reasoning and Acceptance in a Context, *Mind* 101, pp. 1-15
- COHEN, L.J. 1981. Subjective probability and the paradox of the gatecrasher, *Arizona State Law Journal* 2(2)
- COHEN 1995. *An essay on Belief and Acceptance*, OUP
- COHEN L.J. 2002. Should a Jury Say What it Believes or What it Accepts? *In Knowledge and Language. Boston Studies in the Philosophy of Science*, 227, Springer
- DAHLMAN, C. 2017. Determining the Base Rate for Guilt. *Law, Probability and Risk*
- DANCY, J. 1985. *Introduction to Contemporary Epistemology*, Oxford: Blackwell
- DREGER, A. 2015. Galileo's Middle Finger: Heretics, Activists, and One Scholar's Search for Justice, Pinguin
- EBERT P. & SMITH M. (dir.) 2012. *Dialectica*, Special Issue: Outright belief and degree of belief, 66(3)

- ENGEL, P. (dir.) 2000. *Believing and Accepting*. Dordrecht: Kluwer Academic Publishers
- FENTON, N., LAGNADO, D., DAHLMAN, C., NEIL, M. 2017. The opportunity prior: a simple and practical solution to the prior probability problem for legal cases. *ICAIL '17. Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law*, pp. 69-76
- FRIEDMAN, J. 2017. Why Suspend Judging? *Noûs*, 51(2): 302-326
- FRIEDMAN, H. Forthcoming. Inquiry and Belief. *Noûs*.
- HAHN, U., & OAKSFORD, M. 2007, "The Burden of Proof and Its Role in Argumentation", *Argumentation*, 21:39–61
- HARDWIG, J. 1985. Epistemic dependence. *The Journal of Philosophy*, 82(7)
- HAWTHORNE, J., ROTHSCHILD, D., & SPECTRE, L. 2016. Belief is weak. *Philosophical Studies*, 173 (5)
- HUNEMAN, P. 2015, "The True Cost of Conspiracy Theories", *Open Democracy*, <https://www.opendemocracy.net/conspiracy/suspect-science/philippe-huneman/true-cost-of-conspiracy-theories>.
- JEFFREY, R. 1956. Valuation and acceptance of scientific hypotheses. *Philosophy of Science* 23(3): 237-246.
- JOHN, S. 2018, "Epistemic Trust and the Ethics of Science Communication: Against Transparency, Openness, Transparency, and Honesty." *Social Epistemology* 32, 2: 75-87
- KAPLAN J, 1965. Decision Theory and Reasonable Doubt. In Allen, L.E. and Caldwell, M. E. (eds.). *Communication Sciences And The Law: Reflections From The Jurimetrics Conference*, Indianapolis: Bobbs- Merrill Company Inc, 254, 255
- KAPLAN, J. 1968. Decision Theory and the Factfinding Process, 20 *Stanford Law Review* 1065, 1069
- LACKEY, J. 2007. Norms of assertion. *Noûs*, 41
- LAUDAN, L. 2008. *Truth, Error, and Criminal Law: An Essay in Legal Epistemology*. Cambridge Studies in Philosophy and Law
- LAUDAN, L. and SAUNDERS, H. D. 2009. Rethinking the Criminal Standard of Proof: Seeking Consensus about the Utilities of Trial Outcomes, 7 *International Commentary on Evidence* 1, 3, 4
- LEE, M. B. 2018. On doubt. *Philosophia*, 46 (1): 141-158.
- LEITGEB, H. 2013. Reducing belief simpliciter to degrees of belief. *Annals of pure and applied logic*, 164
- LEVI, I. 1960. Must the scientist make value judgments? *The Journal of Philosophy* 57(11): 345-357.
- LILLQUIST, E. 2002. Recasting Reasonable Doubt: Decision Theory and the Virtues of Variability, 36 *UC Davis Law Review* 85, 107–111.

- MCREADY, E. 2014. *Reliability in Pragmatics*. OUP
- MAHER, P. 1993. *Betting on Theories*. Cambridge University Press
- MANDELBAUM, E. 2014. "Thinking is Believing". *Inquiry: An Interdisciplinary Journal of Philosophy*, 57(1): 55-96
- MOON, A. 2018. The nature of doubt and a new puzzle about belief, doubt, and confidence. *Synthese*, 195 (4): 1827-1848.
- ORESQUES, N. & CONWAY, E. 2010, *Merchants of doubt*. Bloomsbury Publ.
- PEIRCE, C.S. 1877. The fixation of belief. *In Popular Science Monthly*, 12, pp. 1-15
- PICINALI, F. 2013. Two Meanings of 'Reasonableness': Dispelling the 'Floating' Reasonable Doubt, *The Modern Law Review* 76(5), pp. 845-875
- POLLOCK, J. 1986. *Contemporary Theories of Knowledge*, Totowa: Rowman and Littlefield
- PRITCHARD, D. 2005. *Epistemic Luck*, Oxford University Press
- ROBERTS, P., & ZUCKERMANN, A. 2010. *Criminal Evidence*, Oxford University Press, 2nd ed.
- ROSS, J. and SCHROEDER, M. 2012, Belief, Credence, and Pragmatic Encroachment. *Philosophy and Phenomenological Research*, 88(2) : 259-288
- SALMON, N. 1995. Being of two minds: Belief with doubt. *Noûs* 29(1)
- SCHUM, D. 1994. *The Evidential Foundations of Probabilistic Reasoning*, Northwestern University Press
- SIMONS, M. 2007. Observations on embedding verbs, evidentiality, and presupposition. *Lingua* 117(6)
- THAGARD, P. (2004). What is doubt and when is it reasonable? *In* M. Ezcurdia, R. Stainton & C. Viger (dir.) *New Essays in the Philosophy of Language and Mind. Canadian Journal of Philosophy, Supplementary Volume*. University of Calgary Press
- WEATHERSON, B. 2005. Can We Do Without Pragmatic Encroachment?. *Philosophical Perspectives*, 19, pp. 417–43
- WEDGWOOD, R. 2012, Outright belief. *Dialectica*, 66(3), 309-329