University of Tartu
Institute of Philosophy and Semiotics

In Defense of Beliefs as Stably High Credences:

How Stability Theory of Belief Avoids the Problem of Conviction based on Purely Statistical Evidence

Master's Thesis in Philosophy

Sophio Machavariani

Supervisor: Alexander Stewart Davies

Tartu, 2021

ეძღვნება ლუკას

Acknowledgments

The thesis wouldn't have been written without the help and encouragement from some people. I would like to express gratitude to Alex Davies for supervising my research and his guidance. I want to thank Javier Belastegui for his comments on the thesis; for being a great teacher and friend. Special thanks should go to the members of the reading group on the stability theory of belief from where I got the idea for this thesis (with no particular order): Tamaz Tokhadze, Simon Mcgregor and Richard Lohse. I am grateful to Patrick Shirreff for being my reviewer and putting great effort to make the thesis better. I also have to thank the participants of the MA seminar at the university of Tartu for their helpful comments on my thesis. Special thanks go to Pyro Suarez for the encouragement I receive from him. I want to thank my mother Eka for standing by me and my partner Luka for his immense support and love.

Table of Contents

Acknowledgments ..... *ii*

Introduction *1*

1. Theories of rational belief and intuition about the legal use of statistics *4*

2. The stability theory of belief (STB) and Staffel's criticism of it *7*

2.1. STB **7**

2.1.1. How STB works 10

2.1.2. The Poss-variant of the Humean Thesis vs the Lockean Thesis 20

2.2. Staffel's criticism of STB **22**

2.2.1. A response from an STB follower: Why is (a) problematic? 24

2.2.2. A response from an STB follower: Why is (b) problematic? 26

3. The defense of STB: how STB avoids the problem of conviction based on purely statistical evidence *29*

3.1. The credence-1 Proposal Fails **30**

3.2. STB and the constraints on narratives for and against guilt **31**

3.2.1 The narrative-based approach to legal fact-finding and the problem of purely statistical evidence 32

3.2.2. STB in the narrative-based framework: explicating the Beyond Reasonable Doubt standard 35

4. STB and the pragmatic constraints on narratives in action: the Simonshaven case *40*

5. STB as a necessary condition on rational belief *46*

5.1. STB is a necessary condition for rationalizing a belief in guilt **46**

5.2. Dismissing the irrelevant coin-flip problem **47**

Summary *50*

List of references *52*

Introduction

Suppose you survive a car accident where you saw that a bus crashed into your car but given that it was nighttime, you could detect neither the color nor the number of the bus which quickly disappeared on the road. You know that there are overall 100 buses in town and that 90% of the buses are operated by the Blue Bus company. The Red Bus company operates 10% of the buses. Based on these statistics you sue the Blue Bus company. Further suppose that the civil standard of proof used is the preponderance of evidence, which means that one ought to find the company liable if it's more likely than not that the company caused the accident. The statistics support the liability of the Blue Bus company with 0.9 probability which does make the liability more likely than not. So, by the standard of proof used, the company should be found liable.[1] Yet, it's counterintuitive to find the company liable based solely on statistical evidence.[2] Since cases like the Blue Bus case can be constructed both for civil and criminal trials, we can generalize and say that it's counterintuitive to believe in guilt or liability based on purely statistical evidence. And, intuitively, any theory of rational belief should be required to accommodate this intuition and disallow a belief in guilt/liability based on statistics alone. Julia Staffel (2016, 2021) argues that a recent theory of rational belief – the stability theory of belief (STB) (Leitgeb, 2013, 2014, 2015, 2017, 2021) – allows for a belief in guilt/liability based on statistics alone, and, hence, fails to provide a sufficient condition on rational belief. Staffel argues that STB cannot provide a necessary condition either since, for STB to disallow a belief based solely on statistics, it has to be paired with pragmatic constraints on rational

[1] One way to deal with the puzzle, is to argue for the modification of legal standards of proof. Particularly people in the literature try to model standards of proof either probabilistically or non-probabilistically so that the legal puzzles are avoided. Backes (2020) dubs this 'the epistemic-to-legal direction of exchange'. This, I'd add, is further branched into 'probabilistic to legal' and 'informal epistemic to legal' directions of exchange. For examples of the probabilistic to legal exachnge, see: Cheng (2012), Kaplow (2014), Urbaniak (2018).

What we're interested in in the paper however is not how epistemic theories of rationality could help us solve legal paradoxes like the Blue Bus case. We're interested in what the paradoxes such as the Blue Bus case should inform our epistemic theories. Backes (2020) calls this 'the legal-to-epistemic direction of exchange'. For examples see Buchak (2014), Smith (2016), Littlejohn (2020).

[2] The Blue Bus puzzle originates in Thomson (1986).

belief. However, once the theory relies on pragmatic constraints, it becomes redundant. Thus, the theory is argued to be neither necessary nor sufficient for rationalizing belief.

In this paper, I defend STB as a statement of a necessary condition on rational belief. The defense of STB is restricted to criminal trials, in particular to cases where the standard of proof used is the beyond reasonable doubt standard. The defense relies on the narrative-based approach to legal fact-finding and the constraints on narratives for and against the guilt hypothesis developed in Di Bello (2013) and Urbaniak (2018). In particular, I propose to pair STB in court contexts with certain pragmatic constraints on the narrative advanced for the guilt hypothesis. Thus complemented STB is later shown to avoid the problem of allowing a belief in guilt based on statistics alone. Moreover, we'll see that thus complemented STB doesn't become redundant and plays an active role in rationalizing a belief in guilt.

So, STB is shown to have resources to avoid the problem of conviction based on purely statistical evidence and it's shown to retain the status of a necessary condition on rational belief in a court context. Thus, Staffel cannot infer from the behavior STB exhibits in court contexts that the theory provides neither a necessary nor a sufficient condition on rational belief. STB is the kind of theory of rational belief that Staffel herself should want in a court context.

In section 1, I introduce the problem of conviction based on purely statistical evidence and the so-called proof-paradoxes. We then consider how this problem might inform theories of rational belief. In section 2, I present the stability theory of belief (dubbed STB) and Staffel's criticism of it based on the alleged incapability of STB to handle the problem of conviction based on purely statistical evidence. In sections 3-5, I embark upon my defense of STB against Staffel's criticism and draw out its legal application.

A note

In what follows, I will be concerned with the Anglo-American legal system for two reasons: this is the system where the standard of proof of interest in this paper – the beyond reasonable doubt standard (BRDS) - is explicitly used. Secondly, the authors I'm engaging with in the paper, implicitly or explicitly, invoke the Anglo-American legal system.

I also want to note that Staffel makes a case against STB in her papers of 2016 and

2021. In the latter, the argumentative strategy slightly diverges from that in the former. I rely mostly on her 2021, however, in 2.2.2. I use considerations against the necessity of STB from her 2016 paper.

1. Theories of rational belief and intuition about the legal use of statistics

In this section, I describe an intuition against the permissibility of believing in guilt/liability (where 'believing in guilt' will be used interchangeably with 'convicting') based solely on statistics and show what this intuition should inform epistemic theories of rational belief. This will set the stage for Staffel's criticism of STB since her criticism is based on the thesis that theories of rational belief should accommodate our intuitions about the use of statistics in a court context.[3] The kinds of theories of rational belief that we'll be interested in bridge the two, standardly accepted kinds of doxastic attitudes: beliefs and degrees of beliefs (i.e. credences). It is a standard assumption in the literature that doxastic attitudes are of two kinds: 1) an all-or-nothing belief which subsumes the categories of belief; disbelief and suspension of judgment, and, 2) degrees of belief (i.e., credences) which afford us a more fine-grained representation of our belief system in the following way. Credence is a real number in the interval of 0 - 1 assigned to a proposition, representing how confident the agent is in the truth of the proposition, where 0 and 1 represent minimal and maximal levels of confidence respectively. We'll be interested in the principles that postulate plausible rationality constraints on how beliefs and credences ought to relate. We'll be interested to find out how well these bridge principles accommodate the intuition against the permissibility of believing in guilt/liability based solely on statistics.

Such intuition is triggered by legal puzzle cases of purely statistical evidence or the so-called proof-paradoxes. Consider the following civil puzzle case:

> **The Blue Bus case** (Thomson, 1986)
> Suppose you survive a car accident where you saw that a bus crashed into your car but given that it was nighttime, you could detect neither the color nor the number of the bus which quickly disappeared on the road. You know that there are overall 100 buses in town and that 90% of the buses are operated by the Blue Bus company. The Red Bus company operates 10% of the buses. Based on these statistics you sue the Blue Bus company. Now, further suppose that the civil standard of proof used is the

[3] For a critical view on this thesis see Backes (2020).

preponderance of the evidence, which means that one ought to find the company liable if it's more likely than not that the company caused the accident. The statistics support the liability of the Blue bus company with 0.9 probability which does make the liability more likely than not. So, by the standard of proof used in civil trials the company should be found liable. Still, it doesn't seem right to find the company liable based on statistics alone.

It doesn't seem right to find the company liable just because it operates 90% of the buses in town. In trying to explain what causes this intuition authors invoke an important feature that statistical evidence arguably lacks, e.g., lack of sensitivity (Enoch et al, 2012); statistical evidence being causally inert (Thomson, 1986); lack of normic support (Smith, 2017)).[4] I won't cover these accounts and in what follows will simply grant the intuition.

Now given this intuition we should expect a good theory bridging rational belief and credence to disallow a belief in guilt based solely on statistics. Staffel (2021) explores whether plausible bridge principles on belief and credence fulfill this expectation.

A plausible rationality constraint bridging beliefs and credences, dubbed the Lockean Thesis in the literature, gives a counterintuitive verdict in the legal puzzles of purely statistical evidence:

The Lockean Thesis: it's rational to believe that P iff one's credence in P is sufficiently high (Following Foley (1993, 2009))

One's credence in the liability of the Blue Bus company is 0.9 which, if any credence, qualifies as 'sufficiently high', and, hence, by the lights of the Lockean Thesis, one ought to believe in the liability of the Blue Bus company. This, though, is counterintuitive: something seems to be lacking for a belief in the liability to be justified.

The Lockean Thesis, arguably, gives counterintuitive verdicts in the puzzles related to the legal puzzle of purely statistical evidence, namely, the standard Lottery Paradox (Kyburg,

[4] These are the features that, according to the authors mentioned, non-statistical evidence has or lacks respectively. This, they think, explains why the non-statistical evidence is a good kind of evidence in contrast with statistical evidence.

1961) and the Harman-Vogel Paradox[5]. Staffel goes through all of these three puzzles and explores whether the modified versions of the Lockean Thesis handle them properly. She considers the stability theory of belief (STB) (Leitgeb, 2013, 2014, 2015, 2017)[6] and what she dubs 'The odds threshold rule' as modified versions of the Lockean Thesis.[7] She argues that these two modified versions of the Lockean Thesis solve the standard Lottery Paradox, but, just like the Lockean Thesis, they both fail at handling the legal puzzle of purely statistical evidence and the Harman-Vogel Paradox.

In what follows I focus the attention of this paper on just STB and the criticism it faces from the legal puzzle of purely statistical evidence. I argue that STB, in contrast with the Lockean Thesis, has resources to handle the legal puzzle of purely statistical evidence and hence is fruitfully deployable in judicial decision-making. Let me briefly say what motivates this choice besides the space limit: I'm not going to address the criticism that STB is unable to handle the Harman-Vogel Paradox because the criticism succeeds but does so in a trivial way. Staffel assumes that the right way to solve the paradox is to disallow a belief that one's lottery ticket is going to lose. She later argues that STB has problems disallowing such a belief. This criticism succeeds in a trivial way because STB cannot be defended from it: for Leitgeb's aim in constructing STB is to allow for a belief in one's lottery ticket losing in a way that's consistent with one's other beliefs.[8] So, criticizing STB on the grounds that it has problems disallowing a belief in one's lottery ticket losing, amounts to saying that STB shouldn't reach the aim it intends to reach. So, Staffel's criticism of STB based on the Harman-Vogel paradox succeeds but in a trivial way.

On the other hand, criticism based on the puzzle of purely statistical evidence provides a novel and a non-trivial challenge to STB, and, as I will show, STB can be defended from it. STB in court contexts can avoid the problem of allowing a belief in guilt/liability based on bare statistics.

---

[5] Based on Nagel (2011), Staffel calls this paradox the Harman–Vogel paradox. It originates in Harman 1973, Vogel 1990.

[6] Similar theory defended in Pedersen and Arló-Costa (2012).

[7] The view dubbed 'The odds threshold rule' is from Levi (1996) and is defended by Lin and Kelly (2012a, 2012b)

[8] This is one way in which Leitgeb solves the lottery paradox (Leitgeb 2017, 2021).

Before I get to present the full version of Staffel's argument against STB, let's look at this theory itself.

2. The stability theory of belief (STB) and Staffel's criticism of it

As we saw in section 1, the Lockean Thesis gives a counterintuitive verdict in the Blue Bus case. It allows for a finding of liability based solely on statistics. Staffel argues that STB is susceptible to the same objection. In this section we'll take a look at STB and examine Staffel's criticism for it. In the presentation of STB the reader can skip the formalism and proceed with a general understanding of the theory.

2.1. STB

Once we accept that we have two distinct ways of understanding our doxastic attitudes, i.e., in terms of beliefs and credences, it becomes legitimate to ask how the two relate to each other. Specifically, how confident should one be that P is true for this level of confidence to qualify as a belief that P. This question invites us to consider plausible bridge principles stipulating how beliefs and credences ought to relate. We already saw one such bridge principle in the form of the Lockean Thesis and we saw it fail to accommodate the intuition triggered by the legal puzzle cases of purely statistical evidence. The stability theory of belief, STB, constitutes another kind of bridge principle which, as I will later show, has resources to accommodate this intuition. Let's have a look at this theory.

In his book Hannes Leitgeb (2017) formulates the following bridge principle on beliefs and credences:

> It is rational to believe a proposition just in case it is rational to assign a stably high degree of belief to it. (Leitgeb, 2017, p42)

As we can see, the principle puts emphasis on the stability property, so let's first expand on what motivates it and ask: why need one to invoke the property of stability in relation to beliefs and credences?

The importance of stability in beliefs has been argued for prior to Leitgeb. As Leitgeb himself notes, Hume is very likely to have held a view on beliefs that rests on the property of stability:

> […] an opinion or belief is nothing but an idea, that is different from a fiction, not in the nature or the order of its parts, but in the manner of its being conceived [. . .] An idea assented to feels different from a fictitious idea, that the fancy alone presents to us: And this different feeling I endeavour to explain by calling it a superior force, or vivacity or solidity, or firmness, or steadiness. (A Treatise of Human Nature, section VII, part III, book I)

The latter properties all seem to invoke the property of stability (or resiliency) of beliefs. An answer to the question about why one would need to invoke such a property is also found in Hume:

> It gives them [i.e., the ideas of the judgement] more force and influence; makes them appear of greater importance; infixes them in the mind; and renders them the governing principles of all our actions. (A Treatise of Human Nature, section VII, part III, book I)

The properties of 'solidity', 'firmness', 'steadiness' are meant to account for the role beliefs play in guiding action; providing a necessary condition for assertion; providing a necessary condition for engaging in suppositional reasoning, etc. To be more specific, the kind of action that is at issue here is extended: an action that takes time to be executed and that requires some, plausibly a high degree of deliberation from an agent who's planning to execute it. For instance, suppose I'm planning to assert what I believe to my friend and am trying to get his agreement. To achieve this, I need to take time and consider the possibilities that could defeat what I believe. In case I still hold the belief in question, I can assert it to my friend and wait for his agreement.

Consider an example: suppose you and I are in a business meeting. It seems to me that P is worth asserting during the meeting, where P denotes 'the current situation on the market speaks in favor of the decision to increase our company's market shares'. I intend to assert P in the hope of getting agreement on your part. But, for my assertion to succeed, I have to make sure that my take on P doesn't change when confronted with the possibilities that I consider at the time. E.g., I should ask myself the following questions: does P retain high credence given that with the increase we'll have to dedicate considerably more time to assessing risks? Or, given that with the increase, we'll have to handle more complaints, possibly even lawsuits from consumers? If, in the face of these potential defeaters, I still think that P, then, P is rationally believable, and I can assert it to you with the hope of getting your agreement.

The aspect that I have to take a moment to check whether my belief is stable in the face of potential defeaters is what makes the action extended (where the action is understood in broad terms: the action of asserting; the action of engaging in suppositional reasoning; or, physically performing some action). And, the stability property of belief seems to account for the role beliefs play in guiding extended action.

One might suggest that the Lockean thesis from section one, and its requirement that credence be sufficiently high, is enough to guide extended action. There is a crucial difference though between high credence and stably high credence, as we'll consider in later sections. To give a brief example, one might have a credence of 1 in the guilt of a defendant which by the Lockean thesis should qualify as one's belief, and, yet, the proposition that the defendant is guilty might not retain credence as high as 1 when its potential defeaters are considered. This difference between high credence and stably high credence is crucial for a proper account of rational belief. For, if a rational belief is just sufficiently high credence, it's not yet guaranteed that this belief will maintain high credence in the face of possible defeaters. Whereas, stably high credence would afford us exactly such a guarantee against potential defeaters.

2.1.1. How STB works

In this section, I present how the stability theory of belief works. So how shall we understand the statement that a belief should be stable in the face of potential defeaters?

A defeater will be understood as a proposition consideration of which could decrease the probability of the proposition the stability of which we're testing. In order to check whether a proposition is stable in the face of potential defeaters, we have to check whether credence in the proposition conditional on each of the potential defeaters stays above some threshold value r, where $0.5 \leq r < 1$. So, stability of belief is understood as credence that remains *stably high enough* under conditionalization on each member of the set of potential defeaters. The 'high enough' is made precise by the threshold r, whereas the 'stably' is made precise by the set of potential defeaters.

Thus, the stability theory works with two parameters: the threshold r, and, the set of potential defeaters conditional on the members of which the stability of a proposition is tested. This is the statement of the core thesis of STB with the relevant two parameters r and Γ:

The Humean Thesis HT^{r}_{Γ}

For all X: Bel(X) iff for all Y, if $Y \in \Gamma$, Y is consistent with X, and, $P(Y) > 0$, then, $P(X|Y) > r$

According to the Humean Thesis X is rationally believable iff for all propositions Y which are members of the set of potential defeaters Γ; which are consistent with X; and, and which have credence assigned strictly above 0, the credence assigned to X conditional on Y is above some threshold value r. If the right-hand side of the biconditional is satisfied, credence assigned to X is stably high and X is rationally believable (i.e., Bel(X)), and the other way round. (In line with Leitgeb, conditionalization will be defined in terms of the ratio formula: $P(X|Y) = P(X \cap Y) / P(Y)$).

Since there are different ways to fill in the two parameters, r, and Γ, we have different versions of the Humean Thesis. Thus, STB already subsumes different plausible bridge principles on belief and credence. Here's an example of how STB can amount to the Lockean Thesis from the previous section.

Suppose the set of potential defeaters Γ is specified as follows: $Y \in \Gamma$ iff $P(Y)=1$. A proposition Y is in the set of potential defeaters Γ iff it has credence 1 assigned, i.e., one is probabilistically certain of Y. Given a thus specified set of potential defeaters, the Humean Thesis from above will amount to the following principle:

Bel(X) iff for all Y, if $P(Y)=1$, then, $P(X|Y) > r$

But this version of the Humean Thesis is just the Lockean Thesis considered in the previous section, since if $P(Y)=1$, then $P(X|Y)$ amounts to $P(X)$. Thus the version of the Humean Thesis from above is the Lockean Thesis:

Bel(X) iff $P(X) > r$

Once we specify r as 'sufficiently high', as required by the Lockean Thesis from section 1, the resulting version of the Humean Thesis where $Y \in \Gamma$ iff $P(Y)=1$ will be just the Lockean Thesis.

As we can see, the Lockean Thesis specifies the set of potential defeaters as the set of propositions that one is probabilistically certain of. One is probabilistically certain of not many things, so the stability that the Lockean Thesis provides to our beliefs is weak: a belief in a proposition will be guaranteed to be stable only over a restricted number of propositions.

The version of the Humean Thesis that Leitgeb focuses on and that Staffel criticizes provides us with a stronger concept of stability, i.e., stability of belief that can be tested over more interesting and not as restricted a set of potential defeaters as the one given by the Lockean Thesis. In line with Leitgeb, I will refer to this version of the Humean Thesis as the Poss-variant of the Humean Thesis:

The Poss-variant of the Humean Thesis[9]

[9] I'm not introducing the notion of P-stability separately but with the Humean Thesis. The notion of P-stability is not explicitly introduced in Leitgeb (2021) either.

Bel(X) iff for all Y, if Poss (Y) (i.e., not Bel(not-Y)), Y is consistent with X, and, $P(Y) > 0$, then, $P(X|Y) > r$

Here Y is a potential defeater iff it is doxastically possible for the agent, i.e., the agent does not believe in the negation of Y, or, entertains Y as possible.[10] In what follows we'll be concerned with the Poss-variant of the Humean Thesis and we'll specify the threshold r as the threshold of 0.5. To put the Poss-variant of the Humean Thesis with the threshold of 0.5 informally:

The Poss-variant of the Humean Thesis informally:
X is rationally believable iff the credence assigned to X conditional on each of the propositions that an agent considers doxastically possible (where the proposition is consistent with X and has credence assigned strictly above 0) is strictly greater than 0.5[11]

As we can see, STB is not a purely formal theory in the sense that the theory cannot issue what's rationally believable based solely on the credence function of an agent. For the theory to issue what's rationally believable, besides the credence function, it needs to take into account which propositions are doxastically possible for an agent, where the latter is determined by pragmatic considerations: by what an agent is interested in; the question that an agent is seeking to answer; the stakes that an agent faces and such. We'll see in the Blue Bus example below how exactly the pragmatic constraints play out.

STB applied to the Blue Bus case

Consider W, a set of possible worlds. Assume a probability distribution P over W.[12] Propositions that are in the probability space W are represented as subsets of W, i.e., $X \subseteq W$

[10] For a response to a concern about circularity involved in the thesis (since 'Bel' occurs on the both sides of the biconditional) see Leitgeb, 2017, p88.

[11] Leitgeb (2017, p121) further proves that one ought to believe in everything entailed by such a proposition.

[12] To be more precise, we assign probabilities to the sigma algebra of W (to the power set of W). See footnote 14 for more details on sigma algebra.

where X is a proposition. A probability distribution over this space is interpreted as credence distribution. Consistency is defined as follows: two propositions are consistent iff they have a non-empty intersection, i.e., $X \cap Y \neq \varnothing$ where X and Y are propositions.

Now consider the Blue Bus case from above except that the total number of buses in town is 10; there is 1 red bus and 9 blue buses. The standard of proof is the preponderance of evidence: find liable iff it's more likely than not that the defendant is liable. Suppose I'm following the Poss-variant of the Humean Thesis and what I want to determine is whether the hypothesis about the liability of the company is rationally believable. Hence, I want to determine whether the credence assigned to the hypothesis that the Blue Bus company is liable remains strictly above 0.5 conditional on each of the propositions that I consider doxastically possible. Since I'm interested in whether the Blue Bus company is liable or not, I consider the following two doxastic possibilities: The proposition that the Blue Bus company is liable, and, the proposition that the Blue Bus company is not liable. So, we cut W into two mutually exclusive and exhaustive possibilities, and, hence, we get a partition Π of W, a set of two sets of possible worlds given below:

Π = {{w: the Blue Bus company is liable in w}, {w : the Blue Bus company is not liable in w}}

{w: the Blue Bus company is liable in w} is the set of all the possible worlds where the company is liable and {w : the Blue Bus company is not liable in w }is the set of all the possible worlds where the company is not liable. Let 'L' and 'N' respectively denote:

L = {w: the Blue Bus company is liable in w}
N= {w: the Blue Bus company is not liable in w}

We can visualize the partition Π in the following way:

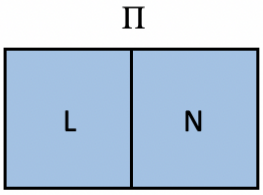

In the given case, an agent partitions the possibility space into two possibilities. But, what exactly determines what possibilities an agent considers? As said, various pragmatic matters do. In this example, we are interested in determining whether the Blue Bus company is liable. This interest prompts the two-celled partition Π.

Now, presumably, a rational agent has a strength of confidence in each of the possibilities she considers. In more precise terms, presumably, she has a credence distribution over these possibilities. Given this, we can think of the partition of the possibility space from above as a credence function (or, a degree of belief function) of an agent. For example, statistical evidence of the Blue Bus case suggests the following coarse-grained credence distribution:

P(L) = P({w: the Blue Bus company is liable in w})=0.9
P(N)= P({w: the Blue Bus company is not liable in w}) =0.1

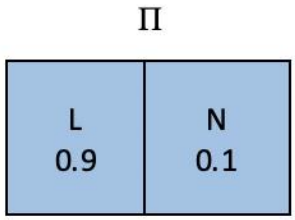

Importantly, one can cut the possibility space as finely or coarsely as one wants. For example, suppose I drop my interest in whether the Blue Bus company is liable, and, instead, embark to answer the question of the form: which one of the buses circulating in town is liable for the accident? Answering this question invites one to restructure the possibility space in a

way that makes the possibility of the liability of each bus explicit. Such a partition would involve cells each one of which represents the proposition that a bus #n is liable. We'll mark this partition of the possibility space as Π'.

Π' = {{w: a blue bus #1 is liable in w}, …, {w : a blue bus #9 is liable in w}, {w : the red bus of the town is liable in w}}

Let 'L1' denote the proposition that blue bus #1 is liable and respectively for the rest of the blue buses up to 'L9' denoting the proposition that blue bus #9 is liable. And, let 'L10' denote the proposition that the red bus of the town is liable.

This fine-grained partition can be visualized in the following way.

Π'

L1 0.1	L2 0.1	L3 0.1	L4 0.1	L5 0.1	L6 0.1	L7 0.1	L8 0.1	L9 0.1	L10 0.1

As we see, each cell gets assigned a credence of 0.1. This is so because the set - {w: the Blue Bus company is liable in w} from the coarse-grained partition Π, now gets cut into further sets of worlds: the sets from {w: a blue bus #1 is liable in w} up to and including {w: a blue bus #9 is liable w}. As a result, the 0.9 credence initially assigned to {w: the Blue Bus company is liable in w} now gets distributed over the more fine-grained, nine sets of worlds.

Now that we've visualized the process of considering possibilities in both coarse and fine-grained ways, let's (on the example of both of these partitions) determine which propositions have stably high credences assigned and hence which propositions are rationally believable according to STB.

Let's start with the coarse-grained partition Π. To determine whether a proposition has stably high credence assigned, we need to check whether its credence conditional on each of the possibilities of the partition is strictly above 0.5. The two sets of possible worlds in Π on p16 give us four propositions, or four possibilities[13]:

{L} = {{w: the Blue Bus company is liable in w}}

{N} = {{w: the Blue Bus company is not liable in w}}

{L} ∪ {N} = {L, N}= {{w: the Blue Bus company is liable in w},{w : the Blue Bus company is not liable in w}}

{L} ∩ {N} = ∅

We get the new probability measure P* (determined by the probability measure P from p17) distributed over the four propositions:

P*({L}) = 0.9

P*({N}) = 0.1

P *({L, N}) = 1

P*({L}∩{N})=P(∅)=0

We should now check whether credence in the relevant proposition is stably high (strictly above 0.5) conditional on each of these four propositions where each of these

[13] Probabilities are real numbers assigned to each of the elements of sigma algebra, where sigma algebra is the closure of the basic set (in this case basic set is Π) under the operations of union, intersection, and complement. So, to be more precise about what we assign probabilities to, we should consider the sigma algebra of Π, i.e. σ Π:
{L} = {{w: the Blue Bus company is liable in w}}
{N} = {{w: the Blue Bus company is not liable in w}}
{L} ∪ {N} = {L, N}= {{w: the Blue Bus company is liable in w},{w : the Blue Bus company is not liable in w}}
{L} ∩ {N} = ∅

We now put in place the new probability measure based on the old one, i.e., based on P. Call this new measure P*. P* assigns probabilities to the subsets of W from above.

This new P* probability measure is determined from the old one (i.e., P). (This is so because the union of the elements of the sigma algebra (in this case of σ Π) will be just Π, i.e., {L, N}). (See Leitgeb, 2017, pp137-39 for a thorough exposition of how a new probability space is formed from the old one).

propositions is compatible with the relevant proposition and has credence assigned strictly above 0.

Let's check whether {L} has stably high credence assigned. The *only* proposition that is compatible with {L} (i.e., has a non-zero intersection with it), and also has credence strictly above 0 assigned, is {L, N}. Thus the only element we can conditionalize {L} on is {L, N}:

$$P(\{L\} \mid \{L, N\}) = P(\{L\} \cap \{L, N\}) / P(\{L, N\}) = P(\{L\}) / P(\{L, N\}) = 0.9 / 1 = 0.9 > 0.5$$ [14]

P({L}|{L ,N}) > 0.5, hence, {L} has stably high credence assigned and is rationally believable. Similarly, we can get that W itself is also rationally believable. {N}, on the other hand, doesn't have stably high credence assigned and is not rationally believable. Consider:

$$P(\{N\} \mid \{L, N\}) = P(\{N\} \cap \{L, N\}) / P(\{L, N\}) = P(\{N\}) / P(\{L, N\}) = 0.1 / 1 = 0.1 \not> ½$$

The only proposition that we could conditionalize {L}, and, {N} on, in this example, is {L, N}. However, a more complex partition than the ones considered in the Blue Bus case, will give more propositions and hence might require us to conditionalize a relevant proposition on many possibilities. Fortunately, Leitgeb provides a simple algorithm with which to identify propositions that have stably high credence assigned (see in the footnote).[15] But, with or without the algorithm, we get that the following are the propositions that have stably high credences assigned in Π:

[14] Where'/' stands for division and where, again, conditionalization is defined in terms of the ratio formula: P(X|Y) = P(X∩Y) / P(Y). The credence in a disjunction is calculated by adding the credences assigned to the disjuncts (assuming that the disjuncts are mutually exclusive).

[15] The algorithm for determining P-stable sets (Leitgeb, 2017, p123): "Assume that $W = \{w1,...,wn\}$, and $P(\{w1\}) \geq P(\{w2\}) \geq \ldots \geq P(\{wn\})$. If $P(\{w1\}) > P(\{w2\}) + \ldots + P(\{wn\})$ then $\{w1\}$ is the first, and least, non-empty *P*-stable set, and one moves on to the list $P(\{w2\}), \ldots, P(\{wn\})$; e.g. if $P(\{w2\}) > P(\{w3\}) + \ldots + P(\{wn\})$, then $\{w1, w2\}$ would be the next *P*-stable set. On the other hand, if $P(\{w1\}) \leq P(\{w2\}) + \ldots + P(\{wn\})$ then consider $P(\{w2\})$: if it is greater than $P(\{w3\}) + \ldots + P(\{wn\})$ then $\{w1, w2\}$ is the first *P*-stable set, and one moves on to the list $P(\{w3\}), \ldots, P(\{wn\})$; but if $P(\{w2\})$ is less than or equal to $P(\{w3\}) + \ldots + P(\{wn\})$ then consider $P(\{w3\})$: and so forth. The procedure is terminated when the least subset of *W* of probability 1 is reached". (See, for the motivation of this algorithm Leitgeb, 2017, Appendix B).

{L, N} t=1

{L} t=0.9

Each of these two propositions is rationally believable. As we see, each of them comes with a threshold t for rational belief. What this means is that if out of these rationally believable propositions you adopt a belief that {L}, then the credence assigned to it (i.e., 0.9) will act as a threshold for rational belief.

Both {L} and {L, N} are rationally believable, but whether one just believes in {L, N} or also adopts a belief in {L} depends on contextual and pragmatic factors. For instance, this can be determined by how bold or cautious one is epistemically: As one moves bottom-up, the more cautious one becomes in what she believes. For, if all one believes in is {L, N}, one makes the most cautious choice since all she believes in now is a tautology. So, bottom-up, a belief becomes weaker. {L}, on the other hand, is the boldest choice here.

In this illustration of STB, we see that STB is not purely formal. When using STB, it's not sufficient to apply the Poss-variant of the Humean Thesis and determine which propositions have stably high credences assigned. One also needs to consider pragmatic constraints at play. We already saw how the pragmatics enter the belief-formation process based on STB. They enter twice: First, for STB to be put to work, an agent has to choose to partition the possibility space in some way. This process is underpinned by pragmatic considerations, e.g., by what an agent is interested in at a time. STB needs some pragmatic constraints at a later stage too: having rationalized which propositions are rationally believable, the theory leaves it up to contextual and pragmatic factors to determine in which one of them to adopt a belief (e.g., whether to believe {L} or not in the Blue Bus case).

Let's now consider how STB applies to the fine-grained partition Π' of the Blue Bus case. The question that underpins the partition Π' is: Which one of the buses in town caused the accident? Here the agent is no longer interested in the liability of the Blue Bus company, the way she was when she considered the partition Π, but in the liability of each bus of the town:

Π' = {{w: a blue bus #1 is liable w}, …, {w : a blue bus #9 is liable w}, {w : the red bus of the town is liable w}}

L1= {w: a blue bus #1 is liable w}

L9 ={w: a blue bus #9 is liable w}
L10 = {w : the red bus of the town is liable w}

10 possible worlds are going to give us 2^{10}=1024 propositions. Thus we have 1024 potential defeaters for each of the 10 propositions. Since would be exhausting to check for the stability of each of the ten propositions, this is where Leitgeb's algorithm comes in handy. However, I will give an example with the standard calculation:

Suppose we want to check whether the proposition 'The Blue Bus company is liable' has stably high credence assigned, i.e., the proposition that one of the 9 blue buses is liable. That is, the union of the first 9 propositions, that is: {L1, L2, L3, L4, L5, L6, L7, L8, L9}. We can check the stability of this proposition relative to {L9, L10}, i.e., the proposition that either a blue bus #9 is liable or the red bus of the town is liable.

P ({L1, …, L9}|{L9, L10}) = P({L1, …, L9} ∩ {L9, L10}) / P({L9, L10}) = P ({L9}) / P ({L9, L10}) = 0.1 / 0.2 = 0. 5 $\not\succ$ 0.5

Given the potential defeater of the form 'Either the blue bus #9 is liable or the red bus is liable', the credence assigned to 'The Blue Bus company is liable' is not stable. Since we found a defeater conditional on which the proposition 'The Blue Bus company is liable' does not have stably high credence assigned, we can infer that the latter proposition is not rationally believable.

The algorithm further shows that the only proposition that is rationally believable in Π' is a long disjunction of the form: {L1, L2, L3, L4, L5, L6, L7, L8, L9, L10}.Thus, all one can believe in Π' is the proposition that one of the buses in town is liable.

If we look back at the intuition against the permissibility of a belief in guilt or liability based solely on statistics and examine how well STB captures it, we'll see that the theory permits a belief in the liability of the Blue Bus company in the coarse-grained partition but leaves it up to pragmatic considerations whether an agent adopts a belief in it or not. Alternatively, if an agent starts instead with considering the fine-grained partition Π', STB doesn't permit a belief in the liability of the Blue Bus company at all. So, the question of whether STB accommodates our legal intuition cannot be answered straightforwardly: It all depends on what pragmatic constraints one will pair STB with.

2.1.2. The Poss-variant of the Humean Thesis vs the Lockean Thesis

I will now consider the importance of the stability property and the difference between the Poss-variant of the Humean thesis and the Lockean thesis with regards to adequately capturing this property. This will make a case for the superiority of the Poss-variant of the Humean thesis to the Lockean Thesis. Thus, this will block a possible concern that the Poss-variant is idle and that the Lockean Thesis and pragmatic constraints can suffice for rationalizing belief.

One of the potential defeaters that we considered when determining the liability of the Blue Bus company was a disjunction of all the propositions that the partition afforded us. That is, we considered as a potential defeater a disjunctive proposition that has credence 1 assigned (e.g., {L, N} in Π). Given this, one might suggest that the Lockean thesis seems to be sufficient for determining rationally believable propositions since all it demands is a conditionalization on the propositions that have a credence of 1 assigned. And, thus, one will conclude that the Poss-variant of the Humean Thesis is nothing over and above the Lockean Thesis.

This reasoning is wrong. The fine-grained partition of the Blue Bus case provides a counterexample to it. Even though in the Poss-variant of the Humean Thesis we'll always have to conditionalize a proposition on a potential defeater of credence 1 (and, hence, use the Lockean Thesis), this will not be sufficient to determine whether a proposition has stably high credence assigned or not. Consider a case from above:

The proposition that one of the blue buses is liable (i.e., {L1, L2, L3, L4, L5, L6, L7, L8, L9}) conditional on the disjunctive propositions that one of the buses is liable (i.e., {L1, L2, L3, L4, L5, L6, L7, L8, L9, L10}) is $0.9/1 = 0.9 > 0.5$. Thus, if all we consider here is the Lockean Thesis and all we conditionalize the former proposition on is a tautology, we'll get that the proposition {L1, L2, L3, L4, L5, L6, L7, L8, L9} has stably high credence assigned and is rationally believable. The Lockean Thesis wouldn't require us to conditionalize it on anything else. And, yet, there are potential defeaters with a credence of less than 1 conditional on which {L1, L2, L3, L4, L5, L6, L7, L8, L9} does not retain stably high credence, and, hence, this proposition is not rationally believable. E.g., such a defeater is {L9, L10}, as we considered above. Thus, conditionalization on the whole partition (and using only the Lockean thesis) does not guarantee that we have determined that a proposition has stably high credence assigned.

Thus, the use of only the Lockean Thesis and checking stability under the propositions that have credence 1 assigned might make us declare propositions as rationally believable where they are not. The Poss-variant of the Humean Thesis captures the notion of stability of belief better and gives a more accurate understanding of it than the Lockean Thesis does.

Yet, the Poss-variant of the Humean Thesis is not as simple as the Lockean Thesis . The latter doesn't require one to perform any kind of conditionalization (or, any kind of conditionalization on propositions that have a credence of less than 1 assigned). The Lockean Thesis only requires a check on whether the credence assigned to a proposition is above some threshold value. This difference in simplicity is innocuous since as said in 2.1.1., STB most adequately governs beliefs that undergird extended action: an action that takes some deliberation to be properly accomplished (e.g., the conviction of a defendant; or the deliberation at the business meeting). Given this profile of STB, it is a good feature of the theory that it requires one to consider potential defeaters at the level of grain that an agent wishes to and, then, take time to check whether a proposition retains stably high credence in the face of these potential defeaters.[16]

I now move on to present Staffel's argument for the thesis that STB fails as a theory and that a belief cannot be stably high credence.

[16] For a more detailed comparison of the Humean and the Lockean theses, where a similar line of thought is present, see Pettigrew (2015).

2.2. Staffel's criticism of STB

Recall from section 1 that a good bridge principle connecting beliefs and credences should be required to accommodate the intuition that it's wrong to believe in the guilt or liability of a defendant based on statistical evidence alone. Staffel argues that STB fails to accommodate this intuition and, hence, rational belief cannot be stably high credence.

In the coarse-grained partition of the Blue Bus case above STB allows for a belief in the liability of the Blue Bus company but leaves it up to the relevant pragmatic constraints whether to adopt a belief in it or not. Staffel though argues that the theory fails to accommodate the legal intuition. This is true if the only part of STB that's considered is the Poss-variant of the Humean Thesis (or any other version of the Humean Thesis): for the Poss-variant of the Humean Thesis indeed determines rationally believable propositions one of which is the proposition that the Blue Bus company is liable. STB however is not just the Poss-variant of the Humean Thesis: a crucial part of the theory is the pragmatic constraints which are to determine the partition of the possibility space, and, whether to believe in a rationally believable proposition or not.

However, the separation between STB, as constituted solely by the Poss-variant of the Humean Thesis on one hand, and, the relevant pragmatic constraints, on the other hand, is an important assumption in Staffel's objection to STB. So, in what follows I will grant that the two are indeed separate and we'll refer to them as 'a pair' or 'a combination'.

Thus we saw earlier that STB, the Poss-variant of the Humean Thesis, makes it rationally permissible to believe in the liability of the Blue Bus company. Staffel argues that for STB to accommodate the legal intuition, it should be paired with one of the following two pragmatic constraints: (a) the one laying restriction on which partition of possibility space should be considered, or, (b) the one laying restriction on which proposition out of the rationally believable propositions to adopt a belief in. She argues that (a) is unworkable. If, an STB follower were to choose the strategy outlined in (b) though, the Poss-variant would be made

redundant by the pragmatic constraints it would have to be paired with. I summarize this attack on STB in the argument below relying on both Staffel's 2016 and 2021.

Staffel's argument against STB

P1 Legal puzzle cases, like that of the Blue Bus case, trigger an intuition that a belief in guilt based on statistical evidence alone is rationally impermissible

P2 Theories of rational belief should be required to accommodate the legal intuition given in P1

P3 STB can accommodate this intuition and disallow a belief in guilt based on statistical evidence alone by introducing pragmatic constraints which either (a) lay restriction on which partition to choose, or, (b) lay restriction on which proposition out of the rationally believable propositions to adopt a belief in.

P4 Given P3, STB does not provide a sufficient condition on rational belief. It has to be complemented with pragmatic constraints

P5 (a) is unworkable since there is a case where the kind of restriction on the partition choice that STB needs is not available.

P6 (b) is unworkable too, since STB when paired with pragmatic constraints becomes redundant

P7 Given P6, STB does not even provide a necessary condition on rational belief

C Given P4 and P7, STB is neither necessary nor sufficient a condition for rational belief and, thus, rational belief cannot be stably high credence

In this thesis, I grant P1 and P2. I also grant that STB, the Poss-variant of the Humean Thesis, is not sufficient for rationalizing belief. It indeed needs to be paired with pragmatic constraints in order to be put to work (as was shown in the Blue Bus example and as is explicit in Leitgeb's framework). What I disagree with is P6. I show that STB can be paired with certain pragmatic constraints in court contexts in such a way that the theory is not redundant and also avoids the problem of purely statistical evidence (thus, capturing the intuition given in P1).

I now present what (a) and (b) responses amount to and Staffel's reasons for thinking why they are not available to an STB follower.

2.2.1. A response from an STB follower: Why is (a) problematic?

As we saw earlier, when we applied STB to the Blue Bus case in the coarse-grained partition Π, STB allowed for a belief in the liability of the Blue Bus company based solely on statistics. Staffel considers a way out on behalf of an STB follower, and, suggests that STB accommodate the intuition by laying restrictions on how the possibility space should be structured, i.e., which partition one should choose to consider.

If an STB follower establishes that a fine-grained partition (Π') is more natural to consider than the coarse-grained one (Π), then, all a rational agent will be permitted to believe is a tautology of the form: one of the buses in town caused the accident and is liable for it. This way STB will avoid the finding of liability of any of the buses based solely on statistics.

Yet, (a) is unworkable because there are cases where it's not possible to have the kind of fine-grained partition where STB ensures a belief in a tautology as the only rationally permissible belief. For STB to ensure a belief in a tautology, the fine-grained partition should have an even credence distribution over it. Yet, there are cases where an even distribution is not available. In such cases, STB seems still to allow for a finding of guilt/liability based on purely statistical evidence.

As an example I will consider a criminal case that's structurally identical to the civil case Staffel presents (so that both civil and criminal cases get covered):

Modified Prisoners[17] (following Staffel (2016, 2021))

Prisoners escape. They come from different sections of the prison and in doing so many of them attack and murder prison guards. Based on the past criminal record of prisoners in each section, we know the proportion of prisoners from each section who would most probably have participated in the killings as opposed to those who would have escaped without killing: 55% of prisoners who escaped from section 1 killed the prison guards; 25% of prisoners who escaped from section 2 killed the prison guards, 15% of prisoners who escaped from section 3 killed the prison guards, 5% of prisoners who escaped from section 4 killed the prison guards. We also know that there are 1000 prisoners in each section and exactly half of the prisoners in each section managed to escape. We also have a list of prisoners from each section.

K1 = {w: a prisoner from section 1 killed the guards in w}
K2 = {w: a prisoner from section 2 killed the guards in w}
K3 = {w: a prisoner from section 3 killed the guards in w}
K4 = {w: a prisoner from section 4 killed the guards in w}

P({K1}) = 0.55, P({K2}) = 0.25, P({K3}) = 0.15, P({K4}) = 0.05,

Now, suppose all of the escapees have been arrested and are brought to court. Only some of them committed murders but we don't know which ones of them did. Staffel claims that since the only evidence available is statistical, one ought not to form an outright belief about the guilt of each defendant. As we'll see, however, STB allows for such a belief.

With the algorithm we get that the following propositions have stably high credences assigned:

{K1, K2, K3, K4}	t=1
{K1, K2, K3}	t=0.95
{K1, K2}	t=0.8
{K1}	t=0.55

[17] The original Prisoners case is from Nesson (1979)

Since {K1} is rationally believable, STB allows for a belief in any of the escapees from section 1 being guilty.

Thus, STB allows for a belief in the guilt of a defendant based on statistical evidence alone. The response based on the restriction on permissible partitions, i.e., (a), does not help an STB follower here. This response is no longer available since, according to Staffel, in the given case one cannot get a fine-grained partition with uniform credence distribution over it.

2.2.2. A response from an STB follower: Why is (b) problematic?

Staffel argues (b) to be problematic for two reasons. Since I will be concerned with each of these reasons in the rest of the thesis, I'll expand on them in different subsections.

2.2.2.1. A credence of 1 proposal

A way an STB follower could accommodate the intuition triggered by legal puzzle cases is by laying restrictions on which proposition out of the rationally believable propositions to adopt a belief in, that is, (b). This though requires that STB be complemented with pragmatic constraints governing which proposition out of the rationally believable propositions one ought to believe in.

What does it mean to lay restrictions on which proposition out of the rationally believable propositions one ought to believe in? For instance, in Modified Prisoners, if an STB follower were able to justify the restriction that out of the propositions that have stably high credences assigned only the tautology be believed (i.e., the proposition of the form: {K1, K2, K3, K4}), then, STB would no longer allow for a conviction based on statistics alone. Yet, to argue for this, one would need to invoke, besides STB, certain pragmatic constraints. Staffel considers a pragmatic constraint of the form:

Credence-1 constraint

> Whenever evidence under consideration is statistical, out of the propositions that according to STB are rationally believable, only the one with a credence of 1 should be believed.

This way, STB would become sensitive to the kind of evidence advanced, i.e., to whether evidence is statistical or not, and would rationalize belief accordingly. However, a problem with this defense of STB is that STB when paired with the pragmatic constraint becomes idle. It is the credence-1 constraint that does the job of avoiding a belief based solely on statistics, not STB (Staffel, 2016). So, we could freely let go of STB and keep the credence-1 constraint, or, else, pair the credence-1 constraint with the Lockean Thesis instead of pairing it with STB (Staffel, 2021, p93). At least, as Staffel argues, there is no reason for why we should keep STB in this picture.

2.2.2.2. The irrelevant coin-flip problem

Later Staffel presents a more general case that purports to undermine the possibility of pairing STB with pragmatic constraints in general (i.e., not just that of the credence-1 constraint from above).

Staffel argues that for STB to preserve the status of a necessary condition on rational belief, it shouldn't be the case that when it's complemented with a plausible pragmatic constraint, STB disallows a belief while the plausible pragmatic constraint allows for it. Unfortunately, there is a case where such a tension between STB and plausible pragmatic constraints arises. In the irrelevant coin-flip problem below, STB doesn't permit a belief in a proposition that a stipulated pragmatic constraint would require one to believe.

The irrelevant coin-flip problem (based on Staffel (2016, 2021), where she follows Fitelson (2015) and Schurz (2019)):

Suppose you have the following credence distribution over the partition consisting of two propositions denoted by Q and not-Q:

P({Q})=0.6 P({not-Q})=0.4

The only rationally believable propositions here, by the lights of STB, are {Q}, and, {Q or not-Q}. Suppose you want to adopt a non-trivial belief in the given partition and hence you adopt a belief in {Q}. Intuitively, now, considering a further proposition, the truth of which is irrelevant to whether Q, shouldn't make you drop a belief in {Q}. For example, considering whether a coin landed heads or tails alongside with {Q} or {not-Q}, shouldn't change your belief that {Q}. This is what the pragmatic constraint is stipulated to require in this example: one shouldn't drop a belief in a proposition when considering propositions that are irrelevant to the former. However, as we'll see, this plausible pragmatic constraint is violated by STB when a more fine-grained partition of possibilities involving {Q} and {not-Q} is considered. Here, {H}is a proposition that heads land, and {T} is a proposition that tails land.

P({Q} and {H})=0.3 P({Q} and {T})=0.3 P({not-Q} and {H})=0.2
P({not-Q} and {T})=0.2

By STB's lights, none of these propositions is rationally believable. Hence, you should abandon your belief that {Q}. But this contradicts what the stipulated pragmatic constraint requires. STB rules out a proposition as not rationally believable when a plausible pragmatic constraint requires one to believe in it, hence, prima facie, we should ignore STB's recommendation and go with what the pragmatic constraint suggests. As a result, STB fails to provide even a necessary condition on rational belief, and, hence, is idle. Thus we get C: STB is neither necessary nor sufficient a condition for rational belief, and, hence, rational belief cannot be stably high credence.

This ends my presentation of Staffel's criticism of STB. In the next section, I provide a defense of STB where we'll see that STB provides a necessary condition on rational belief without succumbing to the problem of conviction based on purely statistical evidence.

3. The defense of STB: how STB avoids the problem of conviction based on purely statistical evidence

This is how my defense of STB will proceed. In 3.1., I argue against the credence-1 proposal. I show why it fails to make a case for the redundancy of STB. In 3.2., I set out to find a pragmatic constraint that STB can be paired with, in court contexts, in such a way that the problem of purely statistical evidence is avoided. The finding of such constraints will, on one hand, make STB accommodate our intuition against the use of bare statistics and will, on the other hand, open up the prospect of using STB in court contexts.

For this purpose, I first cover the so-called narrative-based approach to legal fact-finding which will be shown to preclude the possibility of finding of guilt or liability based on statistical evidence alone. Thus, if a theory can function in the narrative-based framework it will automatically avoid the problem of purely statistical evidence. So, the question I will provide a positive answer to is whether there is a pair of STB and pragmatic constraints that can rationalize a belief beyond reasonable doubt in the narrative-base framework. (As we see, this quest is restricted to criminal trials and the beyond reasonable doubt standard given space limitations, but a similar question can be answered in the positive in the context of civil trials). A plausible explication of the beyond reasonable doubt standard (the BRDS) will reveal a pair of STB and certain pragmatic constraints on accusing narrative that can be fruitfully deployed in narrative-base framework. Hence, we will get a pair that can avoid conviction based on purely statistical evidence. What this shows, again, is that STB works fine in court contexts when the framework for legal fact-finding is narrative-based.

In section 4, I apply the combination of STB and the pragmatic constraints on accusing narrative to a famous criminal case and show how the pair works. In section 5, I show that STB provides a necessary condition on rational belief when paired with the pragmatic constraints on accusing narrative. Here I also dismiss the irrelevant coin-flip problem.

3.1. The credence-1 Proposal Fails

Recall that the credence-1 proposal was meant to complement STB with the pragmatic constraint of the form: whenever the only evidence available is statistical, out of the propositions that according to STB are rationally believable, only the one with credence 1 should be believed. This pairing of STB and the credence-1 constraint was proposed as a way to save STB from the problem of purely statistical evidence. But STB was later thought to be redundant when paired with this constraint.

The credence-1 constraint though fails to make a case for the redundancy of STB since the pairing of STB with this constraint amounts to just another version of the Humean Thesis: the credence-1 constraint requires us to believe in propositions that have credence 1 assigned, i.e., the propositions that we hold to be true no matter what potential defeaters we consider. That is, the propositions that retain stably high credence conditional on any proposition that has credence greater than 0 assigned. This suggests that the credence-1 constraint is just the certainty proposal version of the Humean Thesis:

The certainty proposal
Bel(X) iff P(X) = 1

The certainty proposal is just the following version of the Humean Thesis:

The certainty version of the Humean Thesis
Bel(X) iff for all Y, if P(Y) > 0, Y is consistent with X, then, P(X|Y) > r

Thus, since a pair of STB and the credence-1 constraint is just the certainty proposal version of the Humean Thesis, what the strategy outlined in (b) amounts to is the statement that STB can avoid the problem of purely statistical evidence if what constitutes it is the certainty proposal version of the Humean Thesis. Concern about the redundancy of STB disappears, since the combination of the credence-1 constraint and STB is a special case of STB itself (It's a version of the Humean Thesis).

The discussion here suggests that a proponent of STB can avoid the problem of purely statistical evidence and accommodate the intuition against conviction based on bare statistics by simply invoking the certainty version of the Humean Thesis. This, however, is not an optimal way out of the problem since this version of the Humean Thesis is problematic.

The certainty version of the Humean Thesis allows for a finding of guilt based on statistics alone which is what we wanted to avoid with the adoption of this thesis. Consider a case that illustrates all-or-nothing statistics (Jackson, 2020):

> A case of all-or-nothing statistics (a similar case can be found in Buchak (2014)): Suppose you left your phone in a room with two people: Jim and Jill. You have a clear memory of leaving your phone in the room, yet, when you go back you see it gone while Jim and Jill are still present. You know that 100% of those who steal phones are men. Should you believe that Jim stole your phone?

Given the statistics in the case, the credence assigned to the proposition that Jim stole your phone is 1. Hence, by the certainty version of the Humean thesis, you should believe that Jim stole your phone. This suggests that for a proponent of STB to avoid the problem of purely statistical evidence she should look beyond the certainty version of the Humean Thesis.

This problem with the certainty version of the Humean Thesis also shows that the combination of the Poss variant of the Humean thesis and the credence-1 constraint was doomed to fail in the cases of all-nothing statistics. So, in defense of STB, we should look for a different way out: we should pair the Poss variant of the Humean Thesis with certain other pragmatic constraints where the theory will avoid the problem of purely statistical evidence.

3.2. STB and the constraints on narratives for and against guilt

In what follows, I will propose to pair in court contexts STB (in particular the Poss variant of the Humean Thesis) with certain pragmatic constraints on an accusing narrative. This combination will guarantee that the problem of purely statistical evidence is avoided by STB.

Before I state the proposed combination I present the narrative-based approach to legal fact-finding and show how it precludes the possibility of finding of guilt or liability based on statistical evidence alone. I will then try to answer the question whether there is a pair of STB and pragmatic constraints that can rationalize a belief beyond reasonable doubt in the narrative-base framework, i.e., a pair that can avoid conviction based on purely statistical evidence. Explication of the BRDS will suggest that there is.

3.2.1 The narrative-based approach to legal fact-finding and the problem of purely statistical evidence

Most of the literature on proof paradoxes and the problem of purely statistical evidence misses an important ingredient of judicial adjudication; namely, on what it is that a belief in guilt or liability is, or, should be based on. It is implicitly or explicitly assumed that the belief is, or, should be based on admitted evidence and the claims advanced either in favor or against the defendant. For example, in the Blue Bus case, it is assumed that a belief about liability is, or, should be based on either the claim that statistics support the liability of the company or the claim that it doesn't. It has been argued both at a descriptive and normative level, however, that the finding of guilt or liability is, or, should be based on stories, narratives advanced for and against the defendant which are themselves based on admitted evidence.[18] The normative upshot of why we should require narratives from the defense and the prosecution is that detailed stories about what happened at the crime scene will provide a fair assessment of the culpability of a defendant and will, most probably, help the fact-finder give a just verdict.

Let's think of a narrative for guilt as a set of propositions that forms a coherent story that aims to support the guilt of a defendant. Of course not just about any narrative advanced for guilt will support the guilt hypothesis. There are specific requirements on what the narratives have to be like to support the guilt hypothesis, but I'll consider them slightly later in

[18] See Griffin (2012), Ho (2008) in legal literature; see Di Bello (2013), Urbaniak (2018) in legal philosophy literature; see Pennington and Hastie 1991, 1992, 1993 in psychological literature

this section. In what follows I will use the narrative-based approach to legal fact-finding[19]. I will dub the prosecution's case 'the accusing narrative', and the defense's case 'the defending narrative'.

Now purely statistical evidence either cannot provide a narrative that's good enough to support a conviction, or, it cannot provide a narrative at all. Either way, a conviction will never be issued based on bare statistics in this framework because a narrative constructed on it will be a bad one, or, there will be no narrative constructed at all.

To show the tension between purely statistical evidence and the narratives in court, consider the modified version of Prisoners from section 2.2. All a prosecution can say based on the statistical estimate of what percentage of prisoners in each section could kill prison guards is:

> 55% of the prisoners who escaped from section 1 have a criminal record far worse than the rest of those who escaped from the same section. Since there are a thousand prisoners in section 1 and only exactly half of them managed to escape, we should find 275 of those who escaped guilty of killing the prison guards patrolling section 1. This is so because 275 individuals were very likely to have done it.

These statements clearly do not constitute a narrative. Having said this though, if one were to try to read the statements above as a narrative, one would inevitably conclude that the narrative is not a good one. For example, what the prosecution says suggests that there should be some pieces of evidence available which are not. For instance, the tool with which at least some of 275 defendants could attack the guard; The match between the DNA on the tool and the DNA of at least some of 275 defendants. Yet such evidence is not available in the case.

The case of modified prisoners shows that based on statistical evidence alone one cannot construct a narrative specifying what evolved at the crime scene. Even if we do our best to see the prosecution's case as a story or a narrative in support of the guilt hypothesis it will fail to qualify as a good narrative. What this suggests is that the defendant will never be

[19] For an example of such an approach see Van Koppen and Mackor (2020)

found guilty based on purely statistical evidence since the narrative afforded will fail to qualify as a good one.

But purely statistical evidence might just as well not provide a narrative at all. Consider the Blue Bus case from above where the statistics support the liability of the Blue Bus company with 0.9 probability. It's hard to imagine a narrative about the events of the night of the accident being constructed on it. An absence of narratives about guilt would put the fact-finder in a predicament. Consider: in case there is no narrative presented to a fact-finder, she won't be able to deliberate about conviction. A fact-finder will be in a predicament: since judicial decision-making being strictly binary, a fact-finder has to decide one way or the other (either in favor of the defendant or against her). While a narrative-less setting would make a fact-finder suspend judgment about the guilt hypothesis, and, as a result, paralyze her decision-making.[20] Given this predicament that a fact-finder might face, purely statistical evidence is not something courts would be willing to admit.

Purely statistical evidence will either provide a narrative that is not good or, it won't provide a narrative at all and be dismissed by a fact-finder. In either way, conviction or finding of liability based on bare statistics will be precluded by the framework of narratives. Thus, if a theory can function in this framework it will automatically avoid the problem of purely statistical evidence.

The conflict between the conviction based on a narrative and the conviction based on purely statistical evidence further illustrates where exactly the tension between my defense of STB and Staffel's criticism lies. The background of Staffel's criticism is a plausible requirement on theories of rational belief to accommodate our legal intuitions, in particular the ones prompted by proof paradoxes (like that of the Blue Bus case; or, the modified Prisoners case). If we are to restrict our attention to these idealized proof paradoxes and vindicate the intuition against conviction based on bare statistics, then, STB, as we saw earlier, fails to accommodate the latter intuition. What I suggest though is to rethink what exactly our epistemic theories should accommodate in court context (i.e., revise what exactly the so-called legal-to-epistemic direction of exchange should be about): The discussion in this section implies that epistemic theories should be required to properly handle the cases that are of the

[20] See Ross (2021) for more extended thoughts on judicial decision-making being strictly binary.

right type from a legal perspective, i.e., the cases that involve narratives for and against guilt, and which also happen to preclude finding of guilt based on bare statistics.

As we'll see in the next subsection, a pairing of STB with certain pragmatic constraints can be fruitfully deployed in the narrative-base framework: it can rationalize whether a narrative supports belief in guilt beyond reasonable doubt. Thus, avoiding the problem of purely statistical evidence.

3.2.2. STB in the narrative-based framework: explicating the Beyond Reasonable Doubt standard

I will now show how STB can rationalize whether a narrative supports a belief in guilt beyond reasonable doubt when complemented with certain pragmatic constraints. For this purpose, I first explicate the BRDS itself.

A standard understanding of the BRDS is as follows: one ought to believe in the guilt of a defendant iff it's beyond reasonable doubt that the defendant is guilty. The statement is vague since it doesn't specify what exactly it means for a belief in guilt to be beyond reasonable doubt. Let's give a more precise statement of the BRDS, also the one that's fitting to the narrative-based framework.

> The BRDS
> Believe in the guilt of a defendant iff the accusing narrative supports the guilt hypothesis beyond reasonable doubt, otherwise, acquit

The accusing narrative supports the guilt hypothesis beyond reasonable doubt iff the defending narrative fails to induce reasonable doubt about the guilt of a defendant. But, note that even if the defending narrative fails to induce reasonable doubt about guilt, it might still hold that the accusing narrative doesn't support the guilt hypothesis because it's not strong enough. What this suggests is that a proper statement of the BRDS should involve requirements on both the defending and the accusing narrative. For a belief in guilt to be justified beyond reasonable doubt, the accusing narrative should support the guilt hypothesis

and the defending narrative shouldn't induce a reasonable doubt about the guilt of a defendant. Thus, the right-hand side of the biconditional in the BRDS lumps in two pieces of information that have to be separated.

> The BRDS*
> Believe in the guilt of a defendant iff the accusing narrative supports the guilt hypothesis and the defending narrative doesn't induce reasonable doubt about the guilt of a defendant, otherwise, acquit

A better understanding of the BRDS* hinges on the notion of plausibility as argued convincingly by Di Bello (2020). For consider: in order for a belief in guilt to be beyond reasonable doubt, the narrative in favor of the innocence of a defendant shouldn't induce a reasonable doubt about the guilt of a defendant. In other words, for a belief in guilt to be beyond reasonable doubt, the defending narrative shouldn't be plausible. So, in order to get a more precise understanding of the BRDS*, we need to have a more precise understanding of plausibility when it comes to a defending narrative. Yet, such a precise understanding of the plausibility of the defending narrative is lacking in the literature.[21]

In what follows I will explicate the right-hand side of the biconditional in the BRDS*. I will propose to understand it with the help of the stability theory of belief and certain pragmatic constraints on an accusing narrative.

Let's first consider what it means for the defending narrative to induce reasonable doubt about guilt, i.e., what it means for it to be plausible. I propose to understand the plausibility of the defending narrative in the following way:

> **Plausibility of the defending narrative**
> The defending narrative induces reasonable doubt about the guilt of a defendant (and, hence, is plausible) iff, given the accusing narrative, a fact-finder maintains high

[21] See Di Bello (2020) for an overview of account of this concept of plausibility.

credence in the defending narrative[22]

This requirement on the defending narrative is in line with the requirement STB posits for rational belief, namely that rational belief should remain stable when its potential defeaters (the defeaters that one deems possible) are considered. So, a more precise and elegant understanding of the Plausibility of the defending narrative leads us to STB, in particular to the Poss variant of the Humean thesis.

Plausibility of the defending narrative*

The defending narrative induces reasonable doubt about the guilt of a defendant (and, hence, is plausible) iff a fact-finder maintains a stably high credence in the defending narrative, where a potential defeater is the narrative advanced by the prosecution

STB affords us an elegant formulation of what it means for a defending narrative to induce reasonable doubt about guilt and, hence, be plausible. But this is just an explication of one part of the right-hand side of the biconditional in the BRDS*, namely, that of the defending narrative inducing a reasonable doubt. To get the full explication of the BRDS* we also need to understand what it means for an accusing narrative to support the guilt hypothesis. For this, I draw upon the work of Di Bello (2013) and Urbaniak (2018). Di Bello (2013) states the constraints that narratives for and against the guilt hypothesis have to satisfy for the guilt to be beyond reasonable doubt. Urbaniak (2018) provides a more precise account of these constraints. Following Urbaniak, but skipping formal glosses, I state what the constraints on an accusing narrative should be like.

Explaining Evidence: For a belief in guilt to be justified beyond reasonable doubt it should be the case that the narrative explains the evidence. The narrative explains evidence iff all the pieces of evidence that the narrative doesn't exclude are strongly plausible given the narrative

[22] See Urbaniak (2018) for a different kind of explication to the defending narrative inducing no reasonable doubt (i.e., it being implausible).

Missing evidence 1: For a belief in guilt to be justified beyond reasonable doubt it shouldn't be the case that, given a narrative for guilt, it's strongly plausible that there should be some piece of evidence available and yet it isn't

Missing evidence 2: For a belief in guilt to be justified beyond reasonable doubt it shouldn't be the case that there is a piece of evidence available and yet, given a narrative for guilt, it's strongly plausible that this piece of evidence shouldn't exist

Gaps (or Specificity): For a belief in guilt to be justified beyond reasonable doubt it shouldn't be the case that a certain claim is strongly plausible; it's also strongly plausible that the narrative should contain that claim, and, yet, the narrative doesn't contain it[23]

For the guilt hypothesis to be justified beyond a reasonable doubt, the accusing narrative has to comply with a further requirement. Namely, the narrative forwarded by the prosecution has to survive the challenges presented by the defense. I.e., it has to be resilient to the defense's narrative.

Resiliency: For a belief in guilt to be justified beyond reasonable doubt it should be the case that the accusing narrative is resilient to (1) the narrative(s) advanced by the defense[24], and, (2) to all the evidence that could undermine the accusing narrative

If the accusing narrative satisfies all of these five constraints, we'll take it that it supports the guilt hypothesis. We'll dub it 'a strong accusing narrative'. Thus, having revised the right-hand side of the biconditional in the BRDS* I now state the revised BRDS:

[23] Urbaniak (2018) conceives of these constraints as the criteria by which the prosecution evaluates an accusing narrative relative to the other accusing narratives that he might come up with *pre-trial*. Whereas I'm interested in evaluating narratives for and against guilt from the point of view of a fact-finder, not from the point of view of a prosecutor. So, I will take it that a fact finder considers whether the accusing narrative advanced by the prosecutor complies with the four constraints or not.

[24] "The prosecutor's narrative should resist all challenges coming from the defense lawyer who had an effective opportunity to scrutinize the prosecutor's narrative and who took full advantage of such an opportunity" (Di Bello, 2013, p219)

The revised BRDS:

A belief in the guilt of a defendant beyond reasonable doubt iff the accusing narrative is strong in the defined sense, and, the defending narrative induces no reasonable doubt, i.e., it is not plausible in the defined sense

Here STB afforded us an elegant way of understanding the latter conjunct. It also gives us a neater understanding of one of the pragmatic constraints on the accusing narrative, namely that of Resiliency. The check on the stability of belief is explicitly demanded by the Resiliency constraint of the prosecutor's narrative: what a fact-finder is required to examine here is how confident he remains of the prosecution's narrative after having heard the narrative of the defense lawyer. An elegant way to make the constraint of Resiliency precise is to say that the probability of the accusing narrative should remain stably high conditional on the narrative(s) provided by the defense. The probability assigned to the accusing narrative will be interpreted as the credence a fact-finder is willing to assign to the narrative for guilt. So, checking whether the accusing narrative is resilient will be made precise by checking whether the credence a fact-finder assigns to the accusing narrative is stably high. Since this is the metric by which STB rationalizes belief, STB can most naturally be put to work to decide whether the accusing narrative satisfies Resiliency.

Thus, STB and the pragmatic constraints on accusing narrative can rationalize a belief in guilt as beyond reasonable doubt in the narrative-base framework. But how exactly do these pragmatic constraints that STB is paired with work? Do they lay restriction on what partition to have? Do they constrain which of the rationally believable propositions to adopt a belief in? It is the latter that the pragmatic constraints on narratives help us determine. Consider: What partition one has in a judicial context is fixed by a simple question about which of the narratives advanced by the defense and the prosecution are true. Thus, each cell of a partition, as we'll see in the example in section 4, represents a possible answer to this question, i.e., a narrative advanced either by the defense or the prosecution. Once a partition and a credence distribution over it is set, STB, particularly, the Poss-variant of the Humean Thesis, determines whether the credence assigned to the defense narrative is stably high. If it is, the defense narrative induces reasonable doubt and the conviction should not be due. If it's not though, this still doesn't mean that one ought to convict since the accusing narrative might not

be a strong one. So, one has to assess whether the accusing narrative has the properties captures by the pragmatic constraints. Now, here STB will determine whether the accusing narrative is resilient. If it is, hence if it's rationally believable by the lights of STB, the satisfaction of the rest of the pragmatic constraints will determine whether one should adopt a belief in the guilt of a defendant or not. This is where the pragmatic constraints on accusing narrative play out. If the accusing narrative is not resilient though, then the accusing narrative can be dismissed straightaway, since it already fails to satisfy one of the constraints on accusing narrative (i.e., the resiliency).

As we can see, STB when paired with pragmatic constraints on accusing narrative can be put to work in the narrative-based framework and, hence, can avoid the problem of conviction based on bare statistics. Not to leave the defense of STB at a theoretical level, in section 4, I will apply STB and the pragmatic constraints on accusing narrative to the famous Simonshaven case and will show how they work. Later, in section 5, we'll see that STB in this combination doesn't become redundant.

4. STB and the pragmatic constraints on narratives in action: the Simonshaven case

Let's now consider the famous Simonshaven case and put the combination of STB and the pragmatic constraints on accusing narrative into action. I will skip some of the details of the case. The example will illustrate how this combination allows us to use STB in the courtroom without worrying that the theory will succumb to the problem of purely statistical evidence.

The Simonshaven case (Following Di Bello (2020) and Verheij (2020))

At first instance court: in 2011, in the woods surrounding Simonshaven, Ed Lourens is found in the woods. Nearby, Jenny Lourens is found killed (bloody face, strangled, hit by a metallic object). Cartridges are discovered and gunshot residue is found both on Ed and Jenny. Ed made a call to the emergency with a 40-minute delay. Further, it's

determined that the two separated months ago and Jenny was dating another man which provides the motive of the crime (ED). Ed states that a robber attacked the two from bushes, he fell unconscious and later found Jenny dead (ROBBER). The court also entertains a third possibility, namely that something completely different had happened (D (for different)). The court decided against Ed.

At second instance court: The case was appealed in the second instance court. Here the possibility of a robbery, as claimed by Ed beforehand, was further investigated. It's determined that a man, Perry Sultan, attacked and killed women in the vicinity of the crime scene (PERRY). Yet the involvement of Perry seemed unlikely: the description of the robber by Ed didn't match Perry's; DNA found on the cartridges and the victim didn't match Perry's and the tool with which he committed other crimes did not match with the wounds on Jenny. Still, the possibility that a third party but Perry is the robber is open: after all, the found DNA, the origin of which is unknown, seems to support this option (THIRD). The court dismisses this possibility and also the possibility that something completely different, i.e., 'D', had happened at the crime scene and convicts Ed for good.

As we see, in the first instance, court considered three possibilities as to what could have happened marked with words in capital letters: ED; ROBBER; D. At the second instance, the ROBBER possibility is cut into two further possibilities: PERRY; THIRD.

Following but slightly parting from Verheij (2020), I will conceive of each of these possibilities as a set of propositions; we'll think of these sets of propositions as narratives about what happened at the crime scene; further, these sets of propositions will be thought to be mutually exclusive and exhaustive, and, ordered by a preference ordering. So, we can feel free to represent the narratives as cells of a partition. With preference ordering we determine which narratives are preferred over others and which ones of them are equivalent: A narrative is preferred over another narrative iff the former is compatible with a bigger chunk of evidence than the latter. The preference ordering will be represented by the changing sizes of the cells of a partition. Leaving the representation at that would afford us a purely qualitative representation of narratives, so to make the model more precise, we'll interpret the ordering as

a probability distribution, namely, as a credence distribution:[25]

At first instance court

Fig. 1.

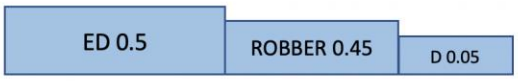

At second instance court

Fig. 2.

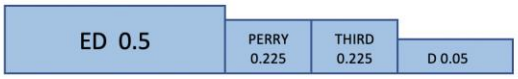

Fig. 3.

As we see at first instance court the case is modeled as a partition of three cells each representing a narrative. Each cell contains only those possible worlds where the given narrative is true. For example, the first cell of the partition in fig. 1., i.e., 'ED' represents all those possible worlds where the narrative labeled 'ED' is true:

ED = {w: the narrative 'ED' is true in w} and respectively for others.

The narrative, as said, is the set of propositions that together imply either guilt or innocence of a defendant, e.g.,: 'ED' is an accusing narrative that could consist of the following sentences: the injuries on Ed and Jenny suggests that the two had a fight; Ed tried to strangle and beat Jenny to death, etc. 'ROBBER' and 'D', on the other hand, are the defending narratives advanced by the defense. As for the 50:50 credence distribution between the

[25] Let me note that Verheij (2020) conceives of the preference ordering as a qualitative relation which can be used to model three different analysis of legal cases: argument-based analysis; scenario (or a narrative) – based analysis and a probabilistic analysis of a case. What I'm doing is to assume from the beginning that the preference ordering is probabilistic and use thus conceived ordering to model the narrative-based analysis of the case with the help of STB and the pragmatic constraints on narratives.

accusing and the defending narratives, it's justified by the presumption of innocence which takes the defendant to be innocent until proven guilty. So, at the beginning of a trial, no matter how incriminating the admitted evidence is, the judgment about the defendant's guilt is suspended. Thus, it's assumed that it's 0.5 probable that the accusing narrative (ED) is true and 0.5 probable that the defending narratives are true (Robber; D) (this assumption is made both at first instance and second instance court). However, a different modeling choice can be made here.

As we saw in the case description, at first instance court Ed was found guilty. Later, at the second instance, 'ROBBER' is extensively investigated and cut into two narratives: 'PERRY' and 'THIRD' (Fig. 2.). 'PERRY' is dismissed by the court, resulting in the partition given in fig. 3. The credence that was assigned to 'PERRY' now gets redistributed to 'THIRD'. The rationale behind this step is that there is evidence that has to be accommodated by the presence of someone else but Ed, namely, the DNA evidence. Since after the exclusion of 'PERRY', 'THIRD' is the most plausible candidate for accommodating this evidence, the credence that was assigned to 'PERRY' now gets redistributed to 'THIRD' (However, a different modeling choice can be made).

Now, on the example of the partition in fig. 3., I will show how the combination of STB and the pragmatic constraints on accusing narrative determines whether the accusing narrative supports guilt beyond a reasonable doubt. Let's start with evaluating the accusing narrative.

Explaining evidence: The accusing narrative in the Simonshaven case, 'ED', explains the evidence. For consider, as defined, the narrative explains evidence iff all those pieces of evidence that the narrative doesn't exclude are strongly plausible given the narrative. The pieces of evidence that 'ED' doesn't exclude, i.e., the gunshot residue on Ed; a strong motive for the crime; the delayed call to the emergency, are indeed strongly plausible given the narrative 'ED'. To put it differently, the presence of these pieces of evidence is indeed explained by the truth of 'ED' and Ed being guilty.

Missing evidence 1: It shouldn't be the case that, given a narrative for guilt, it's strongly plausible that there should be some piece of evidence available and yet it isn't. 'ED' doesn't miss evidence in this sense. 'ED' is constructed based on the three pieces of evidence (the gunshot residue on Ed; a strong motive for the crime; the delayed call to the emergency). So everything that according to 'ED' should be in the evidence, is in the evidence.

Missing evidence 2: It shouldn't be the case that there is a piece of evidence available and yet, given a narrative for guilt, it's strongly plausible that this piece of evidence shouldn't exist. Prima facie, it seems that 'ED' doesn't satisfy this constraint since given 'ED' the DNA evidence of an unknown individual shouldn't exist and yet such evidence exists. This claim assumes that the DNA evidence and 'ED' are inconsistent and for 'ED' to be true the DNA evidence shouldn't exist. However, the two are not inconsistent. For consider, it's very easy for the prosecution to make 'ED' compatible with the finding of DNA by extending 'ED' with a claim about Ed cooperating with another individual whose DNA was found on the cartridges. So, once the accusing narrative is extended with this latter claim, it's no longer susceptible to Missing evidence 2.

Gaps (or Specificity): It shouldn't be the case that a certain claim is strongly plausible, it's also strongly plausible that the accusing narrative should contain that claim, and, yet, the narrative doesn't contain it. E.g., since it's strongly plausible that the victim was beaten by a metallic object, it's also strongly plausible that the accusing narrative should contain information about any link between the object and the defendant, for instance: whether the object carries the DNA of the defendant; whether the defendant was capable of using this object and such. It shouldn't, in principle, be a problem for the prosecution to extend 'ED' so that further claims specifying the link between the metallic object and the defendant are made.

Now as to Resiliency. As we saw, the criteria by which a fact-finder assesses the accusing and defending narratives involves the evaluation of the resiliency of the accusing narrative and the assessment of whether the defending narrative induces reasonable doubt. We also saw that STB provides a precise way in which to determine the resiliency of the accusing narratives and whether the defense narrative induces reasonable doubt. Based on the credence

distribution in fig. 3., STB gets us the following propositions as having stably high credences assigned, and, hence, being rationally believable:

{ED, THIRD, D} t=1
{ED, THIRD} t= 0.9
{ED} t= 0.6

These are the propositions that have stably high credences assigned and each of them is rationally believable. As we see the credence in the accusing narrative 'ED' is stably high. This implies that 'ED' survives the narratives advanced by the defense and, hence, is resilient. On the other hand, we see that neither 'THIRD', nor 'D' have stably high credences assigned. So, neither of them induces reasonable doubt.

Given the pragmatic constraints on narrative, we have established the following: The accusing narrative explains the evidence, doesn't miss evidence in either of the two senses of the Missing evidence, doesn't have gaps, and is resilient. The defending narratives, on the other hand, are shown not to induce reasonable doubt. So, STB and the pragmatic constraints on narratives suggest that the guilt of a defendant is beyond a reasonable doubt. Ed ought to be found guilty, as he indeed was at second instance court. As illustrated in fig. 4., the defense narratives are now eliminated and the probability assigned to them is redistributed to 'ED'.

Fig. 4.

5. STB as a necessary condition on rational belief

In section 4 we put to work the combination of STB and the pragmatic constraints on accusing narrative. We considered a case where this combination helped us determine in the narrative-based framework whether the conviction is beyond reasonable doubt. STB was shown to provide a precise account of whether the accusing narrative involved is resilient and whether the defense narrative induces reasonable doubt. The fact that STB performs well in the cases that involve narratives shows that the theory can avoid the problem of conviction based on bare statistics.

Yet, to fully respond to Staffel's criticism of STB we have to further show that, in this narrative-based framework, STB provides a necessary condition on rational belief and does not become idle when paired with the pragmatic constraints on accusing narratives. This is what I show next. Later, I will dismiss a more general argument undermining the necessity of STB – the irrelevant coin flip problem. This will finish my defense of STB.

5.1. STB is a necessary condition for rationalizing a belief in guilt

Since Staffel might worry that STB becomes idle when paired with pragmatic constraints on accusing narrative, I will here show that the theory retains the status of a necessary condition.

As was proposed, STB informs us what exactly it means for a defense narrative to induce reasonable doubt (or, for it to be plausible) and also gives us a precise understanding of the Resiliency constraint on accusing narrative. Whereas, the four pragmatic constraints on accusing narrative inform us whether it's strong enough to support the guilt hypothesis.

Now if one were to argue that in this proposal STB is idle, one would have to show that at least one of the four pragmatic constraints on accusing narrative captures whether the defense's narrative induces reasonable doubt or whether the resiliency constraint on the accusing narrative is satisfied. It's very unlikely however that any of the four pragmatic constraints on accusing narrative can give a proper account of whether a defense narrative

induces reasonable doubt. This is so because the burden of proof is not symmetrically distributed between the prosecution and the defense: the defense, unlike the prosecution, is not expected to explain why we have a certain piece of evidence available; she's not expected to present a detailed picture of what evolved at the crime scene where the narrative presented doesn't miss evidence in either of the two senses and has no gaps. The defense is only expected to provide a reasonable doubt: make a case in favor of the defendant which seems plausible despite all that the prosecution has argued. Since the four constraints on accusing narrative represent the properties that we do not expect a good defending narrative to have, it's unlikely that any of the four constraints will provide an account of when a defending narrative is good enough (when is it that it induces reasonable doubt, or, is plausible).

This already shows that if the proposal of using STB to explicate the plausibility of the defense narrative is tenable, then, STB is not idle at all. It does a specific job that the pragmatic constraints on accusing narrative cannot do. Thus, STB plays an active part in rationalizing a belief in guilt as beyond reasonable doubt in the narratives-based framework.

5.2. Dismissing the irrelevant coin-flip problem

Still, as we saw in 2.2., Staffel provides an example that purports to undermine a more general claim about STB providing a necessary condition on rational belief. So, let me respond to this criticism.

Recall that Staffel questions the necessity of STB by considering a case where STB rationalizes a belief that is at odds with what a plausible pragmatic constraint would require. This creates a problem for the thesis that STB when paired with pragmatic constraints on rational belief provides a necessary condition on rational belief since this thesis demands that there be no case where STB disallows a belief and a plausible pragmatic constraint allows for it. Here's what the objection from 2.2. looks like when applied to the Simonshaven case (For simplicity, we'll lump the defense narratives 'THIRD' and 'D' under one defense narrative

labeled ‘DEF’).

The credence distribution in the Simonshaven case, a coarse-grained partition:
P({ED})=0.6 P({DEF})=0.4

A more fine-grained credence distribution in the Simonshaven case where for each narrative we consider all the worlds where a coin landed heads (‘H’) and all the worlds where the coin landed tails (‘T’):

P({ED} and {H}) = 0.3
P({ED} and {T}) = 0.3
P({DEF} and {H}) = 0.2
p({DEF} and {T}) = 0.2

In the coarse-grained partition (as we saw already in 3.3.) the narrative ‘ED’ is rationally believable, but not ‘DEF’. The pragmatic constraints are stipulated to require that one does not drop a belief in a proposition when considering a proposition that is irrelevant to the former. Given this constraint, one shouldn’t drop a belief in ‘ED’ when considering propositions about the coin flip. And yet, one does, because none of the propositions from the fine-grained partition from above is rationally believable according to STB. This is the objection from Staffel.

To respond to this problem on behalf of STB, we have to note that the Poss variant of the Humean thesis doesn’t apply across partitions, which is to say, if the thesis declares a belief as stable it doesn’t mean that the stability of this belief is guaranteed across partitions. The application of STB to the Blue Bus case provides an example of this partition-sensitivity: the proposition ‘The blue Bus company is liable’ had stably high credence assigned and was rationally believable in the coarse-grained partition, but it (in the form of ‘One of the buses of the Blue Bus company is liable’) was no longer rationally believable in the fine-grained partition. Now the constraint that an irrelevant proposition shouldn’t affect a belief that P applies across partitions: it demands that ‘ED’, which is rationally believable in the coarse-grained partition from above, be believed in the fine-grained partition of the coin-flip case.

Thus, there is a big asymmetry between STB not applying across partitions and the pragmatic constraint about an irrelevant proposition applying to different partitions. Given this, it's not fair to assess how good the stability theory of belief is by measuring its strength against a constraint that's inherently at odds with how the theory works. If we are fond of the constraint invoked in the coin-flip example and if we want to stick to it, then, probably, we should no longer follow the stability theory of belief. This, however, doesn't show that there's anything wrong with the stability theory of belief since the theory is meant to capture the stability of a belief in a partition, not across partitions.

Another way to respond to the irrelevant coin-flip problem is to say that in using STB's formalism in the coin-flip case Staffel seems to be forgetting what this formalism is intended to be a model of: It's intended to be a model of determining what is rationally believable given a certain partition of the possibility space which is itself determined by certain pragmatic constraints, e.g., by the interests of an agent, by the question he's seeking an answer to, and such. So, STB is not committed to the problem of the coin-flip. STB can be paired with certain other pragmatic constraints, the ones which will not permit the fine-grained partition of the coin-flip example. For instance, if one is interested in considering only the propositions that are relevant to one another, then the fine-grained partition from above will be ruled out for one.

In any case, the irrelevant coin-flip doesn't present a problem for STB in judicial contexts since the partition there should involve the narratives that are relevant to deciding on the guilt of a defendant. It would be strange for a fact-finder to consider propositions of the form "the narrative 'ED' is true and the coin landed heads", as was done in the coin-flip example.

In this section, I provided two reasons for thinking that STB provides a necessary condition on rational belief. One was restricted to court contexts. There we saw that STB when combined with the pragmatic constraints on accusing narrative cannot be let go of since it captures the property that the pragmatic constraints fail to capture, namely that of the plausibility of a defending narrative. Later, I dismissed a more general concern about the necessity of STB which had to do with the irrelevant coin-flip problem.

Summary

In the thesis, I defended STB from Staffel's criticism where she argues that STB provides neither a necessary nor a sufficient condition on rational belief. The reason for this is that STB is argued to allow for a belief in guilt based on purely statistical evidence, which is counterintuitive. I showed that this inference fails: conviction based on a narrative for guilt ensures that the conviction is not issued based on purely statistical evidence. Since in the cases that involve narratives for and against guilt STB (complemented with pragmatic constraints on accusing narrative) performs fine, the theory avoids the problem of purely statistical evidence.

In the end, I dismissed a concern that STB when complemented with pragmatic constraints fails to provide a necessary condition on rational belief.

A corollary of my defense of STB from Staffel's criticism is an account of how STB can be fruitfully deployed in a court context. This makes room for further research on whether the combination of STB and the pragmatic constraints on accusing narratives provide a promising explication of the beyond reasonable doubt standard and, hence, a novel candidate to modeling legal standards of proof in both probabilistic and non-probabilistic ways.

List of references

Backes, M. (2020). Epistemology and the law: why there is no epistemic mileage in legal cases, Philosophical Studies, 177 (9): 2759-2778.

Buchak, L. (2014). Belief, credence, and norms, Philosophical Studies, 169: 285–311.

Cheng, E.K. (2013). Reconceptualizing the burden of proof, Yale LJ, 122:1254

Di Bello, M. (2013) Statistics and probability in criminal trials: The good, the bad and the ugly. Dissertation, Stanford University.

Di Bello, M. (2020). Plausibility and Reasonable Doubt in the Simonshaven Case, Topics in Cognitive Science, 12 (4):1200-1204.

Enoch, D., Spectre, L and Fisher, T. (2012). Statistical Evidence, Sensitivity, and the Legal Value of Knowledge, Philosophy and Public Affairs, 40 (3):197-224.

Foley, R. (1992). Working Without a Net: A Study of Egocentric Epistemology, Oxford University Press, New York.

Foley, R. (2009). Beliefs, degrees of belief, and the Lockean thesis, in F. Huber, and C. Schmidt-Petri (eds), Degrees of belief, Springer, Dordrecht, Synthese Library (Studies in Epistemology, Logic, Methodology, and Philosophy of Science), vol 342.

Griffin, L.K. (2013). Narrative, truth, and trial, Georgetown Law Journal, 101:281–335.

Harman, G. (1973). Thought, Princeton University Press, Princeton.

Ho, H. L. (2008). A Philosophy of Evidence Law: Justice in the Search for Truth, Oxford University Press.

Hume, D. (1978). Treatise of human nature (L. A. Selby-Bigge, Ed.; 2nd ed.). Oxford University Press.

Jackson, E. (2020). Belief, Credence, and Evidence, *Synthese*, 197 (11):5073-5092.

Kaplow, L. (2014). Likelihood ratio tests and legal decision rules, American Law and Economics Review, 16(1):1–39

Kyburg, H.E. (1961). Probability and the Logic of Rational Belief, Wesleyan University Press, Middletown, Connecticut.

Leitgeb, H. (2013). Reducing belief Simpliciter to degrees of belief, Annals of Pure and Applied Logic, 164 (12):1338-1389.

Leitgeb, H. (2014). The Stability Theory of Belief, Philosophical Review, 123 (2):131-171.

Leitgeb, H. (2015). I—The humean thesis on belief, Aristotelian Society Supplementary, 89 (1):143-185.

Leitgeb, H. (2017). The Stability of Belief. How Rational Belief Coheres with Probability, Oxford University Press, Oxford.

Leitgeb, H. (2021). Stability and the Lottery Paradox, in I. Douven (Ed), Lotteries, Knowledge, and Rational Belief: Essays on the Lottery Paradox, Cambridge University, Cambridge, pp. 147-170.

Levi, I. (1991). The fixation of belief and its undoing: Changing beliefs through

inquiry, Cambridge University Press, Cambridge.

Lin, H and Kelly, K (2012a). A geo-logical solution to the lottery paradox, with applications to conditional logic, Synthese, 186(2):531-575.

Lin, H and Kelly, K (2012b). Propositional Reasoning that Tracks Probabilistic Reasoning, Journal of Philosophical Logic, 41 (6):957-981.

Littlejohn, C. (2020). Truth, knowledge, and the standard of proof in criminal law, Synthese, 197 (12):5253-5286.

Nagel, J. (2011). The Psychological Basis of the Harman-Vogel Paradox, Philosophers' Imprint, 11:1-28.

Nesson, C.R. (1979). Reasonable doubt and permissive inferences: The value of complexity, Harvard Law Review, 92(6):1187–1225.

Pedersen A. P. and Arló-Costa. H. (2012). Belief and probability: A general theory of probability cores, International Journal of Approximate Reasoning, Volume 53, Issue 3.

Pennington, N and Hastie, R. (1991). A cognitive theory of juror decision making: the story model, Cardozo Law Review, 13:519–557.

Pennington, N and Hastie, R. (1992). Explaining the evidence: tests of the story model for juror decision making, Journal of Personality and Social Psychology, 62(2):189–206

Pennington, N and Hastie, R. (1993). The story model for juror decision making, Cambridge University Press, Cambridge.

Pettigrew, R. (2015). II—Pluralism about Belief States, Aristotelian Society Supplementary Volume, 89 (1):187-204.

Ross, L. (2021). Rehabilitating Statistical Evidence, Philosophy and Phenomenological Research, 102 (1):3-23.

Schurz, G. (2019). Impossibility Results for Rational Belief, Noûs, 53 (1):134-159.

Smith, M. (2010). What else justification could be, Noûs, 44(1): 10-31

Smith, M. (2016). Between probability and certainty: What justifies belief, Oxford University Press, Oxford.

Staffel, J. (2016). Beliefs, buses and lotteries: Why rational belief can't be stably high credence, Philosophical Studies, 173 (7):1721-1734.

Staffel, J. (2021). Three Puzzles about Lotteries, in I. Douven (ed), Lotteries, Knowledge, and Rational Belief: Essays on the Lottery Paradox, Cambridge University Press, Cambridge, pp. 74-94.

Thomson, J.J. (1986). Liability and individualized evidence, Law and Contemporary Problems, 49(3):199–219.

Urbaniak, R. (2018). Narration in judiciary fact-finding: a probabilistic explication, Artificial Intelligence and Law, 26 (4):345-376.

Van Koppen, P. J. and Mackor, A. R. (2020). A scenario approach to the Simonshaven case, Topics in Cognitive Science, 12 (4):1132-1151.

Verheij, B. (2020). Analyzing the Simonshaven Case With and Without Probabilities, Topics in Cognitive Science, 12 (4):1175-1199.

Vogel, J. (1990). Are there Counterexamples to the Closure Principle, in M. D. Roth

and G. Ross (eds), Doubting: Contemporary Perspectives on Skepiticism, Kluwer Academic Publishers, Dordrecht, pp. 13-29.

Non-exclusive licence to reproduce thesis and make thesis public

I___Sophio Machavariani

1. herewith grant the University of Tartu a free permit (non-exclusive licence) to

reproduce, for the purpose of preservation, including for adding to the DSpace digital archives until the expiry of the term of copyright,

"In Defense of Beliefs as Stably High Credences: How Stability Theory of Belief Avoids the Problem of Conviction based on Purely Statistical Evidence"_____,

supervised by Alexander Stewart Davies _____.

2. I grant the University of Tartu a permit to make the work specified in p. 1 available to the public via the web environment of the University of Tartu, including via the DSpace digital archives, under the Creative Commons licence CC BY NC ND 3.0, which allows, by giving appropriate credit to the author, to reproduce, distribute the work and communicate it to the public, and prohibits the creation of derivative works and any commercial use of the work until the expiry of the term of copyright.

3. I am aware of the fact that the author retains the rights specified in p. 1 and 2.

4. I certify that granting the non-exclusive licence does not infringe other persons' intellectual property rights or rights arising from the personal data protection legislation.

author's name Sophio Machavariani
15/05/2021