Lotteries and the Roads to Knowledge Failure

Anaid Ochoa Echeverria

Department of Philosophy

McGill University, Montreal

June, 2024

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of PhD in Philosophy

Committee members:

Prof. Michael Blome-Tillmann (Supervisor)

Prof. Stephanie Leary (Supervisor)

Prof. Eran Tal (Internal Examiner)

Prof. Ulf Hlobil (External Examiner)

Table of Contents

Abstract / Résumé

There is a wide consensus among epistemologists that we fail to know *lottery* propositions (that is, highly likely propositions solely supported by statistical evidence), but there is no agreed-upon explanation of why we fail to know them. Yet, paradoxes surrounding lotteries continue pressing on the need to identify the correct explanation. This dissertation evaluates various explanations of why we do not know lottery propositions, which normally take one of two forms. The first argues that lottery beliefs are true by mere *epistemic luck* (roughly, luck, chance, or accident in possessing a true belief in a proposition *p*, which prevents us from knowing *p*), and provides an account of epistemic luck that entails that lottery beliefs are epistemically luckily true. Here I evaluate and reject an explanation that relies on the largely dominant account of epistemically lucky beliefs as "unsafe true beliefs". The second argues that lottery beliefs are unjustified (where being justified in believing *p* is necessary for knowing *p*), and provides an account of justification with a condition that lottery beliefs or propositions do not satisfy. I divide conditions on justification into "probabilistic" and "non-probabilistic" and argue, contrary to Douven and Williamson (2016), that justification *can* have a probabilistic condition that lottery beliefs fail to satisfy, thus restoring such accounts as a theoretical possibility to solve the "lottery paradox". I also argue that among the dominant candidates of non-probabilistic conditions, Smith's (2016) normic support condition on justification is the only suitable one to explain why we do not know lottery propositions. Alternatively, if lottery propositions are justified, there is nonetheless a solution to the lottery paradox compatible with their possessing such epistemic status, which consists in rejecting the "aggregativity of justification". I argue that this principle is incorrect, and that preserving it obstructs solving what I identify as the most basic form of the lottery paradox. Finally, I address the legal correlates of lottery propositions –i.e., litigated claims supported by bare statistical

evidence. I divide explanations of why such claims do not meet a given standard of proof into two types: externalist and internalist. I present independent problems for two externalist explanations –Pritchard's (2018, 2022) modal explanation and Blome-Tillmann's (2017) knowledge-first explanation, and argue that they face such problems in virtue of being externalist. This finding motivates the need to evaluate externalist accounts *qua* externalist to determine if an externalist account is a viable solution to the "proof paradox".

………..

Il existe un large consensus en épistémologie sur le fait que nous ne connaissons pas les propositions de *loterie* (c'est-à-dire des propositions hautement probables mais uniquement étayées par des preuves statistiques), mais il n'existe aucune explication convenue à cet égard. Pourtant, les paradoxes entourant les loteries continuent de souligner la nécessité d'identifier la bonne explication. Cette thèse évalue diverses explications expliquant pourquoi nous ne connaissons pas les propositions de loterie, qui prennent normalement l'une des deux formes suivantes. La première soutient que les croyances de loterie sont vraies par simple *chance épistémique* (en gros, la chance, le hasard ou l'accident de posséder une vraie croyance en une proposition *p*, ce qui nous empêche de connaître *p*), et fournit une explication de la chance épistémique qui implique que les croyances de loterie sont vraies par chance épistémique. Ici, j'évalue et rejette une explication fondée sur l'explication largement dominante selon laquelle les croyances épistémiquement chanceuses sont des « vraies croyances non-sécuritaires ». La deuxième soutient que les croyances de loterie sont injustifiées (où être justifié de croire *p* est nécessaire pour connaître *p*), et fournit une explication de la justification avec une condition que les croyances ou propositions de loterie ne satisfont pas. Je divise les conditions de justification en « probabilistes » et « non-probabilistes » et je soutiens, contrairement à Douven et Williamson (2016), que la justification *peut* avoir une condition probabiliste que les croyances de loterie ne satisfont pas, rétablissant ainsi de telles explications en tant que possibilité théorique de résoudre le « paradoxe de la loterie ». Je soutiens également que parmi les candidats dominants aux conditions non probabilistes, la condition de support normique sur la justification de Smith (2016) est la seule appropriée pour expliquer pourquoi nous ne connaissons pas les propositions de loterie. Alternativement, si les propositions de loterie sont justifiées, il existe néanmoins une solution au paradoxe de la loterie compatible avec leur statut épistémique, qui consiste à rejeter « l'agrégativité de la justification ». Je soutiens que ce principe est incorrect et que le préserver empêche de résoudre ce que j'identifie comme la forme la plus fondamentale du paradoxe de la loterie. Finalement, j'aborde les corrélats juridiques des propositions de loterie – c'est-à-dire les réclamations litigieuses étayées par de simples preuves statistiques. Je divise les explications de pourquoi de telles affirmations ne satisfont pas à un critère de preuve donnée en deux types : externalistes et internalistes. Je présente des problèmes indépendants pour deux explications externalistes – l'explication modale de Pritchard (2018, 2022) et l'explication de Blome-Tillmann (2017) fondée sur la connaissance comme concept primitif, et je soutiens qu'elles sont confrontées à de tels problèmes en raison de leur caractère externaliste. Cette découverte justifie la nécessité d'évaluer les explications externalistes en tant qu'externalistes afin de déterminer si une explication externaliste est une solution viable au « paradoxe de la preuve ».

Acknowledgements

I would like to first express the deepest appreciation to my supervisors, Prof. Michael Blome-Tillmann and Prof. Stephanie Leary. Prof. Blome-Tillmann has guided me throughout my doctoral trajectory. I have had the opportunity to engage in continuous discussions with him and have immensely benefited from all of them, even more so from those in which our contrasting opinions and his persistent guidance came together to further push me to grow philosophically. Prof. Blome-Tillmann made these interactions especially gratifying given his sustained disposition to discuss our ideas in a very direct, frank, and open way, while providing the safe space to do so. Prof. Leary has guided me throughout most of my PhD years, and our discussions have always been marked by her stimulating feedback and the philosophical insight gained through them. I always looked forward to our meetings knowing they guaranteed to substantially boost the quality of my thought. I cherish every one of them and I am deeply grateful to Prof. Leary for her time and dedication.

Secondly, I am especially thankful to my dissertation examiners, Prof. Ulf Hlobil and Prof. Eran Tal. Their meticulous and insightful comments not only clarified strengths and shortcomings of parts of my dissertation, but also gave me much to reflect upon and develop in subsequent papers.

Thirdly, I am indebted to the members of the Philosophy Department at McGill. My participation in seminars, teaching assistantships, and discussion groups at the department broadened my philosophical interests substantially and made me a more robust and open thinker. Special thanks to Prof. Ian Gold, Prof. Chris Howard, Dr. Moritz Bodner, Dr. Andre Martin, Dr. Brandon Smith, and Samuel Carlsson Tjernström, for their feedback on parts of my doctoral work and their overall academic support.

I am also indebted to the Epistemology Work-in-Progress Seminar at UNAM, coordinated by Prof. Lourdes Valdivia. During the writing stage of my thesis, I had the opportunity to discuss most of my thesis manuscripts at this seminar. As always, I received incisive and in-depth feedback from faculty and PhD fellows. Special thanks to Prof. Peter Klein, Prof. Lourdes Valdivia, and Dr. Melahuac Hernández, for their invaluable support at this stage of my doctoral studies.

My research would not have been possible without the financial support from the CONACyT doctoral grant I received for four consecutive years, as well as the doctoral fellowship for international students from the Fonds de Recherche, Nature, et Technologies of Quebec for three consecutive years. Finally, I also thank the philosophy department at McGill for granting me the Graduate Excellent Award for six months in order to complete my doctoral dissertation.

1. Introduction

This dissertation is concerned with the epistemic status of beliefs in lottery propositions – that is, highly likely propositions which are solely supported by statistical evidence. Consider the following example: You have bought a ticket in a fair lottery of 10,000 tickets with only one winner. Based on the information that there is a .9999 probability that your ticket loses, you come to believe that your ticket will lose. The proposition that your ticket –say, ticket number 54– is among the losing tickets, is a lottery proposition. Hence, given your evidence, it is extremely likely that your belief that your ticket is a loser is true. As a result, would you say that you *know* that your ticket is a loser? If the answer is 'no', would you then say you are at least *justified in believing* so?

The epistemic status of beliefs in lottery propositions is of interest to epistemology given how conceptually puzzling it is. Consider the lottery proposition that ticket number 54 will lose. Intuitively, despite having bought the ticket knowing its extremely low chances of winning, you do not know that it will lose.[1] After all, you do not throw away the ticket until you learn the lottery results announced through the proper channels. Yet, there are countless propositions about our ordinary world that we take to know based on evidence that does not make them as extremely likely as probabilistic evidence makes lottery propositions. For example, you may come to know that a convenient store is open based on a neighbor's testimony, that your colleague's name is Peter based on your memory of having heard his name when you were introduced, or that the Musée des Beaux-Arts is four blocks away from here based on a local Montrealer's instructions. The chances of the neighbor's and local Montrealer's testimonies being incorrect, or of your

[1] There is wide philosophical agreement that lottery propositions do not constitute knowledge. As will be apparent, all accounts of knowledge and justification discussed in this dissertation share such view.

memory failing in the above circumstances, may be higher than the odds of winning the lottery. Thus, the probability of your beliefs that the store is open, that your colleague's name is Peter, and that the museum is four blocks away from here, may be lower than the probability of your belief that your ticket will lose (given the evidence for each belief).[2] Because we know so many propositions based on evidence that makes them less likely to be true than the evidence for lottery propositions, the intuitive claim that we do not know lottery propositions becomes deeply puzzling. Thus, an explanation for why we fail to know lottery propositions while we know innumerable other less likely claims (given our evidence) is called for. Without such an explanation, skeptical worries loom for a large portion of our ordinary knowledge.

The puzzle above has come to be known as "The Lottery Problem", initially discussed by Gilbert Harman (1968) and continuously addressed up to the present day. The second puzzle dominating the contemporary literature on lottery propositions is known as "The Lottery Paradox", originally presented by Henry Kyburg (1961). Consider again the lottery example. Even though you may not know that your ticket is a loser based merely on the odds, these odds are excellent evidence to believe so – it is .9999 likely that it loses! At face value, such evidence supports your lottery belief enough to (at least) justify it. Notice that, just as is the case with your lottery ticket, for any other lottery ticket *i*, there is a corresponding lottery proposition of the form ⟨ticket #*i* will lose⟩, whose probability of being true given the evidence is also .9999. Therefore, for every lottery proposition, you are justified in believing that it is true.[3] Aggregating your justified propositions

---

[2] We might even come to *know* the same proposition that your ticket is a loser based on evidence that renders it *less* likely to be true than mere probabilistic evidence renders it, such as by hearing the lottery results in the radio. Given how large the lottery is, it may be that the probability of the radio broadcaster reading the digits incorrectly or having received a misprinted copy of the lottery results is higher than the probability of winning the lottery.

[3] A couple of clarifications are in place. Firstly, for the purposes of this dissertation and unless otherwise specified, I will use 'believing *p*' and 'believing that *p* is true' interchangeably. For a difference in both terms even when accepting the equivalence of '*p*' and '*p* is true', see Kvanvig (2003:34). Secondly, the lottery paradox I am concerned with is about *propositional* justification. Having a justified belief in a lottery proposition (that is, having *doxastic* justification for our lottery beliefs) may be distinguished from being justified in believing or having justification to believe a lottery

such that if you are justified in believing each of them you are justified in believing their conjunction, you can deduce that you are justified in believing that every lottery proposition is true. However, you *know* on the basis of the evidence that not every lottery proposition is true (since you know that one ticket will win). And, assuming knowledge entails justification, you are justified in believing that not every lottery proposition is true. Thus, you seem to be justified in believing both that every lottery proposition is true and that not every lottery proposition is true.

Solving Harman's lottery problem and Kyburg's lottery paradox can serve as a criterion to choose among alternative accounts of knowledge or justification, if not as a condition of adequacy for such accounts.[4] Both puzzles motivate the inquiry for a satisfactory account of why we fail to know lottery propositions. Such an account will identify a condition of knowledge that known propositions about our ordinary world satisfy but lottery propositions do not. If such a condition is justification, an account for why lottery propositions fail to be justified provides an answer to the lottery paradox and may explain away our intuition that we are justified in believing such propositions.

The interest on the status of lottery propositions extends beyond theoretical epistemology and the epistemic standards of knowledge and justification; in particular, it extends over discussions focused on the concepts of legal evidence and legal proof.[5] For example, it is common to conceive standards of legal proof as standards for how much the evidence admitted in court "proves" or supports a litigated claim, where this support is understood solely in terms of how

proposition (that is, *propositional* justification), where the latter is compatible with our not believing the proposition in question (see Silva & Oliveira, 2024).You might not have come to believe every lottery proposition due to fortuitous circumstances – for example, you haven't had the time to entertain the possibility that every lottery proposition is justified for you– and yet still be justified in believing every lottery proposition (once you entertained such possibility, you'd acknowledge you have excellent reason to believe every lottery proposition).

[4] For example, John Hawthorne (2004) and Duncan Pritchard (2008a) considered lottery beliefs to be justified, true beliefs that are not knowledge. According to them, lottery beliefs constitute counterexamples to the traditional analysis of knowledge.

[5] See, for example, H.L. Ho (2021) and M. DiBello (2013).

probable such claim is given the evidence. Understanding standards of proof in probabilistic terms gives rise to a puzzle mirroring the lottery problem, known as "The Proof Paradox". The puzzle contrasts two types of scenarios in which the same litigated claim is supported by different forms of evidence. In the first, the claim is highly supported by bare statistical evidence, e.g., a base rate making it .9 likely. In the second, it is supported by eyewitness positive identification, whose reliability is roughly .7.[6] Even though the claim is more likely given the base rate than given the eyewitness identification, courts would normally not hold someone liable based solely on base rates, while they would do so based on eyewitness identification. And this course of action seems right to us; but why? If legal proof or support is understood only in terms of how probable a litigated claim is given the evidence, bare statistical evidence should be more than sufficient to satisfy some standards of proof (e.g., the preponderance of the evidence standard, stating that the evidence should support claim C more than not-C). Yet, we do not think this type of evidence suffices. Cases involving only this type of evidence have thus come to be known as "legal lottery cases."[7]

In the present dissertation I evaluate dominant accounts on the status of lottery propositions in relation to the epistemic standards of knowledge and justification, with the primary objective of finding appropriate explanations of why we fail to know such propositions. I also discuss the status of litigated claims supported by bare statistical evidence –the legal correlates of lottery propositions– in relation to standards of proof. I complete the dissertation in four main projects.

There are two main types of explanations of why we do not know lottery propositions: those according to which lottery beliefs fail to be justified, and those according to which lottery beliefs fail to meet a condition different from justification, where the only plausible candidate is

[6] See K.L. Fields (2013:1799).
[7] See Pritchard (2018).

the condition of not being "epistemically luckily true." [8] The first project (Section 2) of this dissertation consists in evaluating an explanation of why we do not know lottery propositions insofar as they are true by mere epistemic luck. I evaluate and reject an explanation that relies on the largely dominant account of epistemically lucky beliefs as "unsafe true beliefs".[9] To do so, I raise counterexamples to accounts of epistemic luck in terms of safety conditions that involve both "locally" and "globally" reliable ways of forming beliefs in nearby worlds. All such counterexamples present a lottery case illustrating the next possibility: the process of selecting the lottery winner might be such that any world in which it delivers a different outcome is extremely far away from the actual world. In addition to being a case of "safe ignorance", this type of lottery case shows that, ultimately, either (veritic) epistemic luck is not unsafe true belief or beliefs in lottery propositions are not epistemically luckily true.

The second project (Section 3) begins by examining some of the main views proposed in the literature on the justificatory status of lottery propositions that explain, or suggest an explanation of, why beliefs in such propositions are not justified. Each view is evaluated with respect to whether it is suitable to account for our failure to know these propositions. Some of these views on justification, here called "probabilistic views", claim that lottery propositions do not meet a condition on justification in which a proposition instantiates a *structural* property – roughly, a property fully expressible by means of logical and probabilistic notions.[10] I argue that a problem raised for some structural properties by Douven and Williamson (2006) fails to arise for other similar probabilistic views. This result in conjunction with arguments provided in Section 4

---

[8] Roughly, a belief is epistemically luckily true when, even if justified, it turns out true as a matter of mere accident or coincidence. Reasonably, the fact that a belief turns out true by mere accident prevents it from constituting knowledge. For a general overview of the notion of epistemic luck, see Engel Jr. (2024).

[9] Pritchard is the main proponent of the identification of epistemic (veritic) luck with unsafe true belief. See Pritchard (2005), (2007), (2008a) and (2008b).

[10] See Igor Douven (2002) and Sharon Ryan (1996) for examples of probabilistic views.

allow to conclude, contrary to what Douven and Williamson argue, that justification *can* have a probabilistic condition that lottery beliefs fail to satisfy. Probabilistic accounts of justification are thus restored as a theoretical possibility to solve the lottery paradox (by rejecting that lottery propositions are justified). Finally, I also evaluate four non-probabilistic views that account for why lottery beliefs are not justified: those according to which lottery propositions fail to meet a pragmatic condition (Douven, 2012), a virtue-epistemology condition (Kelp, 2017b), a knowledge-first condition (Ichikawa, 2014), and a "normic support" condition (Smith 2016). I argue that the normic support condition is the best among the four candidates.

The third project (Section 4) consists in addressing the lottery paradox. I present a solution to it by providing two arguments against one of its underlying principles, the *aggregativity* of justification –that is, that justification is a property that a conjunction has if all of its conjuncts have it. This project is done against the backdrop of a proof by Douven and Williamson (2006) criticizing probabilistic views on justification, which (as mentioned above) concludes that lottery propositions do not instantiate a structural property, and crucially, which also assumes the aggregativity of justification. Now, the first argument shows that the aggregativity of justification is incorrect, as well as a restriction of it that is used in Douven and Williamson's proof. The second argument shows that maintaining aggregativity of justification prevents us from solving what I here identify as the lottery paradox in its "most basic form" (where the paradox involving the notion of justification is a "meta-level paradox"), while solving this basic form of the paradox can be easily achieved if such principle is rejected. Both arguments provide a solid basis for best understanding Douven and Williamson's proof as a reductio of the aggregativity of justification, rather than as a reductio of propositions having a structural property (as they intended). Additionally, since the aggregativity of justification is important to assess the justificatory status of conjunctions based on the justificatory status of their conjuncts, I propose and defend an alternative to such principle that avoids the lottery paradox.

Finally, I complete the fourth project in section 5. Many non-probabilistic accounts of legal standards of proof have been put forward in the literature to solve the proof paradox. In this section, I divide all accounts into what I here call 'externalist' and 'internalist' accounts. Externalism about standards of proof establishes a condition on a standard mostly independent of the quality of the admitted evidence for the litigated claim, and explains how claims supported by bare statistical evidence do not satisfy such condition. I provide two case studies of dominant accounts of proof standards, Pritchard's (2018, 2022) "modal view" and Blome-Tillmann's (2017) "knowledge-first probabilistic view" which present an attractive explanation of the insufficiency of bare statistical data to satisfy a standard of proof. I argue that they are both externalist, and that they both independently face serious problems. Afterwards, I argue that their common externalist feature explains why they independently fail. I then conclude by suggesting an internalist desideratum on accounts of standards of proof.

2. Safety, Lotteries, and Failures of the Imagination

1. Introduction

A large variety of safety conditions have been offered in the literature as complete or partial accounts of the nature of knowledge. Roughly, such conditions state that a subject *S*'s belief in a proposition *p* is *safe* when *S* could not easily have falsely believed *p*.[11] In this paper, I restrict the discussion to safety conditions proposed as complete accounts of *epistemic luck*, which aim to explain why certain true and intuitively justified beliefs fail to be knowledge in terms of such beliefs falling prey to a modal veritic type of luck.[12] In particular, I assess how well safety conditions can be used to explain why beliefs in lottery propositions –that is, highly likely propositions which are solely supported by statistical evidence– are epistemically luckily true. The general explanation shared by safety accounts is that even though a proposition such as 〈ticket #34 will lose〉 is extremely likely given its high odds of being true (in this case, that it is .9999 probable that ticket #34 loses), such a proposition could still easily have been false and thus, whether a belief in it turns out true is simply a matter of luck.

The main rationale for treating lottery beliefs as epistemically luckily true is that there appears to be no other explanation of the fact that they fail to be knowledge despite being true, since they seem justified or rationally acceptable on the basis of purely statistical reasoning, given

[11] Ernest Sosa's (1999) and Duncan Pritchard's (2007, 2008, 2012, 2015) accounts of safety as a condition of knowledge are two of the prominent accounts. A third one is Timothy Williamson's (2000), under which safety is a condition of knowledge but not constitutive of it. For a general overview of safety conditions of knowledge, see Danin Rabinowitz (2011).

[12] Roughly, a belief is epistemically luckily true when, even if justified, it turns out true as a matter of mere accident or coincidence. Reasonably, the fact that a belief turns out true by mere accident prevents it from constituting knowledge. For a general overview of the notion of epistemic luck, see Engel Jr. (2010). The main proponent of the identification of epistemic veritic luck with unsafe true belief is Duncan Pritchard (2005, 2007, 2008a and 2009).

their extremely high probability (see Hawthorne (2004) and Pritchard (2007)). Additionally, dominant explanations in terms of lottery propositions lacking justification or rational acceptability face serious problems yet to overcome.[13] Thus, an account of how lottery propositions are epistemically lucky is called for. Further, it seems natural to understand such luck in terms of safety, since there is large agreement that safety is among the best candidates to account for epistemic luck and it has garnered support by evolving to (presumably) successfully account for our failure to know lottery propositions.[14]

This paper pushes back against the previous considerations by showing that there is a type of case involving lottery propositions that inevitably lies beyond the scope of any reasonable safety account of epistemic luck. To achieve this aim, I first elaborate a counterexample to accounts of epistemic luck in terms of safety conditions that involve "locally" reliable ways of forming beliefs (roughly, ways of belief formation that produce more true beliefs in *p* than false beliefs in *p* in close worlds). The counterexample involves a lottery case, dubbed "The Intergalactic Lottery Case", illustrating the next possibility: the process of selecting the lottery winner might be such that any world in which it delivers a different outcome is extremely far away from the actual world.

---

[13] The main views in the literature that explain (or suggest an explanation of) why lottery beliefs are not rationally acceptable propose either what we'll call a "formal" property of rationally-acceptable beliefs or a "non-formal" property, where a formal property is (roughly) a property reducible to logical and probabilistic notions (for examples, see Sharon Ryan (1996), Igor Douven (2002), and Lin & Kelly (2012)). Douven and Williamson (2006) have argued that formal views are untenable by providing a proof to show that if there is a property of rationally acceptable propositions that lottery propositions lack, it cannot be a formal property (although if rational acceptability is relativized to certain formal contexts, some formal properties of rationally-acceptable beliefs so relativized could escape their criticism. For such alternative, see Hannes Leitgeb's (2014)). Salient views that propose non-formal properties in a knowledge-first vein, such as Kelp's (2017) virtue epistemology and Ichikawa's (2014) knowledge-first views of justification, are non-informative with respect to why we fail to know lottery propositions. Another salient view, Douven's pragmatic view (2012), relies on problematic claims about the individuation of the content of beliefs (see section 3.4.1 of this dissertation). A promising property may be captured by Martin Smith's (2016) normic support condition (although see Michael Blome-Tillmann (2020b)).

[14] Safety remained an exception among the various post-Gettier accounts of epistemic luck insofar as it continuously developed into versions that avoided shortcomings arisen for its earlier versions. For example, a version of safety evolved from requiring the target belief to be true in "most close worlds" into a version requiring such belief to be true *also* in "all very close" worlds (Pritchard, 2008b), driven by the need to explain lottery beliefs as unsafe beliefs.

After defending such counterexample from an objection, I present a variation of it against accounts of epistemic luck in terms of safety conditions that involve “globally” reliable ways of forming beliefs (roughly, ways of belief formation that produce more true than false beliefs in close worlds, not only in *p* but in any other proposition produced in the same way). Inspired in safety conditions that involve globally reliable ways of belief formation, I propose an account of epistemic luck in terms of safety that avoids the second counterexample. I argue, however, that with enough imagination, a counterexample to such account may be found. The Intergalactic Lottery Case, in all of its proposed variants, is not only intended to be a case of safe ignorance, but it also shows that either veritic epistemic luck is not unsafe true belief or beliefs in lottery propositions are not true by mere epistemic luck.

2. Safety, Epistemic Luck and Lottery Propositions

To cover all the safety conditions on offer in the literature without examining their various details and peculiarities, we can follow Di Yang (2019) in identifying two types of sub-conditions of any safety condition:

i) a "modal distance" condition, capturing a degree of a relevant type of dissimilarity[15] among the actual world[16] and other possible worlds, in terms of which a special class C of close worlds is defined, and

ii) a "modal frequency" condition, capturing the frequency in which a given event *e* (proposition, fact) obtains in the C-worlds.[17]

A lucky event is then conceived by combining (i) and (ii), as follows:

An event *e* is lucky if, and only if, *e* obtains but the frequency of C-worlds in which *e* obtains is not sufficiently high.

According to this modal conception of luck, beliefs that are epistemically luckily true are dealt with by treating them as unsafe, that is:

> S's belief in *p* is epistemically luckily true if, and only if, S's belief in *p* is true but *unsafe*, that is: S's belief in *p* is true but the frequency of C-worlds in which a belief

---

[15] The relevant type of similarity in safety theories is the relation of overall similarity (similarity considering all respects, rather than similarity in one respect). An exception among safety accounts is Blome-Tillmann's (2020a) "non-reductive safety", which takes modal distance to capture a degree of *epistemic* similarity, where similarity is epistemically restricted to similarity "in those respects that are relevant for knowledge". It is worth exploring whether non-reductive safety can account for why we fail to know the modally stable lottery propositions presented in this paper as part of the counterexamples to safety. However, to cover most safety theories, I will here restrict my analysis to those involving the notion of overall similarity. This could be problematic for the counterexamples here presented if safety conditions are better cashed out with a more restricted notion of similarity rendering such counterexamples ineffective (while counting intuitive cases of knowledge as safe beliefs). Such a restricted notion of similarity, however, would need to be specified before it poses a real problem to the proposed counterexamples. Thanks to Michael Blome-Tillmann and Ulf Hlobil for discussion on this point.

[16] Throughout the paper, I use 'the actual world' to designate the world of evaluation of the relevant safety-claims corresponding to a given scenario (even if the scenario is counterfactual, such as the Intergalactic Lottery Case), rather than our actual world. This sense of "actuality" can be roughly understood in the following way: a world *w* is the actual world in a context *c* if, and only if, it is being evaluated in *c* whether a given modal claim (e.g., a safety-claim) is true or not in *w*.

[17] These sub-conditions are not independent from each other, since the special class C in (ii) is a function of there being a determinate threshold for the dissimilarity relation in (i). Also notice that the class C is defined by the dissimilarity relation *with respect to the actual world*, so C might be a different class when different worlds are taken as the actual world, that is, as the world of evaluation.

of S in *p* is formed in the same way as in the actual world and *p* is true is not sufficiently high.

The formulation of this principle can be made more precise in multiple respects, e.g., by specifying what class C is (the class of *nearby* worlds, of *very nearby* worlds), by specifying what counts as a "way" in which a belief is formed (e.g., a belief-forming method, a belief-forming process, a basis for believing), or by specifying how high the frequency of C-worlds in which a believed proposition is true should be in order for such proposition to count as safe (e.g., it must be true in all C-worlds, most C-worlds, at least one C-world). The lack of specificity of this principle is useful for it cover a large range of safety conditions differing in those respects. Acceptable specifications are constrained by whether the resulting version of the principle covers all and only cases of epistemic luck. Those specifications that count very far off worlds from the actual one as C-worlds, and tolerate only a small number of C-worlds in which a belief is false for it to count as safe, risk counting non-luckily true beliefs as unsafe. For example, if a belief is taken to be safe iff it is true in *all* worlds similar to the actual one in *some* respect, most ordinary beliefs would be unsafe, since they are false in some sceptical world which at least shares some feature with the actual world. Those specifications that only admit very close worlds to the actual one as C-worlds and tolerate a large number of C-worlds in which a belief is false for it to count as safe, risk counting luckily true beliefs as safe (since very close worlds might fail to differ enough from the actual world to make a luckily true belief false). For example, if a belief is taken to be safe iff it is true in some world similar to the actual one in all respects, any true belief would count as safe, including the luckily true ones. In particular, a specification identifying safety with truth in a large number of close worlds clashes with the verdict that beliefs in lottery propositions are unsafe, since they are typically true in a large class of close worlds (given the low chances for lottery propositions to be false, they would be true in the great majority of close worlds). For that reason, throughout this paper I assume that the correct frequency for a belief in a lottery proposition to be unsafe is for such proposition to be false in at least one C-world in which a belief in it is formed in the same way as in the actual world, that is:

Safety:

> S's belief in *p* is epistemically luckily true if, and only if, S's belief in *p* is true but *unsafe*, that is: S's belief in *p* is true but there is some C-world in which a belief of S in *p* is formed in the same way as in the actual world, and *p* is false.

In this paper, I take Safety as adequately capturing most safety conditions intended to explain lottery cases as involving epistemic luck. Safety, however, is still inadequate in that it only encompasses safety conditions whose application to a belief in a given proposition depends on the truth value (with respect to the C-worlds) of that very same proposition. These safety conditions are said to involve a *local* type of reliability[18] because if a belief in a proposition *p* is safe in the sense defined by them, the way in which such belief was formed is supposed to "track" (throughout the C-worlds) the truth of *p* but not the truth of any other proposition (even if such a proposition is the object of a belief formed in the same way). This contrasts with safety conditions involving a *global* type of reliability, where if a belief in a proposition *p* is safe in the sense defined by them, the way in which the belief was formed is supposed to "track" (with respect to C-worlds) the truth of *any* proposition that is the object of a belief formed in the same way. This gives rise to the next variation of Safety:

Global Safety:

[18] See, for example, Dani Rabinowitz' The Safety Condition for Knowledge (2021).

> S's belief in *p* is epistemically luckily true if, and only if, S's belief in *p* is *globally unsafe*, that is: S's belief in *p* is true but there is a C-world in which a belief of S in a proposition *q* is formed in the same way in which S's belief in *p* was formed in the actual world, and *q* is false (where *q* might be different a proposition than *p*).

For the moment, let us leave Global Safety aside. Whichever is the best way to formulate Safety, its application to deal with lottery propositions is very simple. No matter how unlikely it is that a lottery proposition is false, only a small difference in the actual world is required for that to happen. Just a small difference in the lottery winner selection process —e.g., a couple of different numbered balls coming out from a tombola— are needed to yield a different lottery result. Given this, for every lottery proposition, there is a C-world in which a false belief in it is formed in the same way as in the actual world, so the belief is unsafe. Thus, Safety provides a simple explanation of why beliefs in true lottery propositions are epistemically luckily true.

Only two things are needed for Safety to account for the case of lottery propositions: that the modal distance condition does not track the likelihood of the target proposition (so it is not true that the more likely a proposition is, the closer a world in which it is true is to the actual world, and vice versa), and that the differences needed from the actual world to make the target proposition false are small enough to stay below the threshold of modal distance for being a C-world. Unfortunately, the simplicity with which Safety explains lottery cases turns against it, or so I argue next.

3. The Intergalactic Lottery Case: A Counterexample to Safety

Given the simplicity with which Safety deals with beliefs in lottery propositions, the only thing needed to show that it fails is to provide a lottery case in which the process of selecting the winner of the lottery delivers different results from the actual one only in *very far away* possible worlds (given a reasonable specification of the modal distance condition), while maintaining the intuition that the corresponding lottery proposition is not known. Here I provide one such case, having the following structural features:

A) The process by means of which the winning lottery ticket is selected is so spatially large, causally complex, and having a long duration, that a different result could not easily have obtained. Because of the nature of this process, the more a particular result deviates from the actual result, the harder it is that it could have obtained.

B) If (A), then there is no C-world in which a different winning ticket is selected, so any true belief in a relevant lottery proposition is safe (and, therefore, not epistemically luckily true, according to Safety). This is especially true with respect to results differing the most from the actual result.

A scenario satisfying both conditions (A) and (B), while maintaining the intuition that the relevant lottery proposition is not known, constitutes a counterexample to Safety, since the conjunction of both conditions is incompatible with such principle. One such scenario is provided by the following story:

> *Intergalactic Lottery Case*: Every five thousand years, the Intergalactic Federation of Planets (IFP) runs an Intergalactic Lottery in which one (and only one) citizen of the IFP

is prized with the unique opportunity to live a very long and happy life on Planet Paradise IFP (excluding those citizens already living in Planet Paradise). The lottery is run this way because it takes five thousand years for an inhabitant of Planet Paradise to peacefully die, and only then a vacancy for a new inhabitant opens.

The process for selecting the winning ticket is as follows. Firstly, the number of each lottery ticket has as many digits as there are planets in the IFP, and given that the number of planets in the IFP is astronomically large, the number of digits in the number of each lottery ticket is astronomically large as well. Secondly, each planet of the IFP gets to select one digit of the number of the winning ticket by means of some local subprocess with the following characteristics:

a) it is spatially very large, relative to the size of the planet in which it occurs,

b) it is causally very complex, and

c) it has a very long duration (nearly five thousand years).

Given that each of the subprocesses are gigantic at a planetary scale and there is an astronomical number of them, the process for selecting the number of the winning ticket is astronomically larger than each of its gigantic planetary-scale subprocesses.

To illustrate the nature of the subprocesses, consider the following instance (let us keep in mind that, as previously stipulated, any digit-selecting process is comparable in scale, causal complexity and duration to the following one): Planet Fungus creates multiple gigantic enclosed platforms (each one of 16,000 km^2) in which fungus is freely allowed to grow at a very slow rate, with the expectation that a gigantic digit-like pattern naturally forms in one of them. If a platform is filled with fungus and no digit-like pattern forms, the process is restarted in that platform. The first recognizable digit-like pattern to form is the selected digit. Once the digit is obtained, it is securely registered, codified and sent to the organizers of the Intergalactic Lottery.

In this fanciful scenario, an Earth citizen of the IFP, Annie, reasons in the following usual way: "The Intergalactic Lottery has an astronomical number of tickets, so the chances of my ticket being the winner are close to zero. Ergo, if I can be sure of something, it's that my ticket is not the winner." As the result of this reasoning, Annie forms the belief that her ticket is not the winner. And Annie is right, her ticket is not the winner. Thus, Annie ends up forming a true belief in a lottery proposition.

Intuitively, Annie is justified in thinking that her lottery ticket is not the winner (on the basis of purely statistical reasoning), and the proposition believed by her is true, yet it is also intuitive that she does not know such proposition. With respect to these set of intuitions, Annie's belief is no different from any other typical belief in a lottery proposition (except for the lottery proposition that is about the actual winning ticket). However, besides its science-fictional features, the Intergalactic Lottery stands apart from normal lotteries in a crucial aspect: the *astronomical scale and complexity* of the winner-selecting process. To understand why this is crucial, let us restrict our attention momentarily to the digit-selection subprocesses. Given their nature, the counterfactual variations in the outcome of *any* of them need a significant dissimilarity from the actual world, significantly larger than it is required for counterfactual variations in the outcome of any normal digit-selection process (e.g., one of the numbered balls coming out from a tombola is a 6 rather than a 2). Thus, for a *single* digit of the number of the winning ticket to have been different, a significant degree of dissimilarity from the actual world is necessary—e.g., significantly different causal chains of fungus growth during a very long period of time, forming a very different gigantic digit-like pattern than the one produced in the actual world. Thus, we obtain the following intuitive result:

Distant Difference:

> The number of the Intergalactic Lottery's winning ticket could only have differed in one of its digits in far-off worlds

Let us say that *p* is an intergalactic proposition if, and only if, *p* truly reports that a particular ticket of the Intergalactic Lottery is a loser. Accepting Distant Difference entails that, for every intergalactic proposition *p*, *p* is only false in a far-off world, so (contrary to Safety) every typical belief in *p* is safe. Yet, it is intuitive that Annie does not know that her ticket is a loser, despite her belief being true, justified and (according to Distant Difference and Safety) safe. Thus, we have a class of counterexamples to Safety.

It might be argued that Distant Difference is not a decisive threat to Safety. For all we have shown, it might be argued, Distant Difference is compatible with there being a safety condition with a modal distance threshold that counts worlds in which intergalactic propositions are false as C-worlds, according to such safety condition. This is not a rebuttal of the counterexample unless the safety condition in question is specified. However, in the absence of an argument showing that no such safety condition exists, its mere possibility renders the counterexample inconclusive.

Granting for the moment the previous response (there are good reasons to think that it should not be granted, which I discuss shortly), there still are intergalactic propositions for which it does not work. For consider the number of a ticket, *k*, that differs from the number of the actual winner in *all its digits*. And consider an intergalactic proposition p_k reporting that the ticket #*k* is not the winner. A world in which p_k is false (that is, in which the ticket #*k* is the winner) is *astronomically far-off* from the actual world. This is because i) if a number differs in all of its digits from the number of the actual winner, the outcomes of *all the subprocesses* of the Intergalactic Lottery would have to be different for that ticket to be the winner, and ii) a world that differs from the actual world in the outcomes of all such subprocesses is *astronomically far-off* from it (since it involves an astronomical number of planetary-scale differences from the actual world). No specification of the class of C-worlds seems able to count this kind of (astronomical) difference among worlds as below the threshold of a reasonable modal distance condition for Safety, since we expect of such condition to count significantly smaller differences among worlds as above that threshold. To see this, consider the next three variations of the Intergalactic Lottery Case:

> S_1) Annie has the Intergalactic Lottery ticket with #*k*. Given the astronomical number of digits of each number of an Intergalactic Lottery ticket, people need to rely on an extremely powerful and reliable software designed by the IFP (used in extremely powerful and reliable computers) to check and compare the numbers of lottery tickets. Annie forms the belief in the proposition, *p*, that her ticket is not the winner (expressed to herself by the sentence 'My ticket is not the winner') in the way specified above.[19] Based on that belief, she infers that the ticket with *that* number (that is, *k*) is not the winner, whichever is the number of her ticket. Assuming Annie's inference allowed her to form a belief, Annie ends up believing that the ticket #*k* is not the winner, that is: she ends up believing p_k (which, as

[19] I am assuming that Annie formed her belief in *p*, and her belief in p_k, after the result of each subprocess was settled (maybe by years, decades or centuries), so if the counterfactual capturing the relevant safety condition requires that everything before the target belief is formed is kept fixed throughout the possible scenarios required for the evaluation of that counterfactual, that by itself guarantees that there is no nearby world in which Annie formed the belief in *p* (or in p_k) and her ticket was the winner of the Intergalactic Lottery.

It is also interesting to notice that if a Millian view on indexical expressions such as 'my ticket' (that is, the claim that the only semantic content of those expressions is their designatum) is a correct view, 'My ticket is not the winner' expresses the same proposition in the present context as 'Ticket #k is not the winner', regardless of whether Annie is aware that 'my ticket' and 'ticket #k' are co-designative terms in such context. In that case, *p* and p_k are the very same proposition, so the inference form *p* to p_k in S_1 is unnecessary.

we have seen, is true because *k* differs from the number of the winning ticket in each one of its digits).

Based on her previous beliefs, Annie decides to run a scam. Being a very skillful programmer, she creates multiple copies of her lottery ticket that are indistinguishable from the original for the ticket-reading software, although the copies would easily be revealed as fake if they were examined by agents of the IFP with a more sophisticated software. Her plan is to sell each copy to a number of her acquaintances who do not know each other and who live far away from one another. Annie thinks that the scam is likely to go unnoticed if her ticket does not win, so the risk of being discovered is insignificant. After making all the copies for the scam, Annie puts them in a box, together with a list of the people she intends to scam, inside a secret drawer under the desk in her office, and the real lottery ticket next to that box. Annie ends up justifiably and truly believing the proposition, *q*, that her ticket is next to the box.

S_2) Everything is like S_1, except for the following: Annie's boyfriend, Philipp, discovers Annie's secret drawer by accident, and all its contents. Philipp discovers that he is one of the people that Annie intends to scam. Believing that she is going to win the lottery, and to teach her a lesson, Philipp replaces the original ticket with one of the copies and leaves it next to the box, so *q* ends up being false.

S_3) Everything is like in S_1, except for all the facts about the winning ticket selection process that need to be different for *p* and p_k to be false.

Intuitively, the belief in *q* is knowledge in S_1, while it is not knowledge in S_2 (given that *q* is false in S_2). Also, intuitively, the beliefs in *p* and p_k are not knowledge in any of the previous scenarios, even though they are true and (presumably) justified in S_1 and S_2. Let us compare, however, the modal distance between all such scenarios to assess whether these intuitions can be captured by a safety condition.

Let d_1 be the degree of modal distance between S_1 and S_2 (that is, the degree of dissimilarity between S_1 and S_2), and let d_2 be the degree of modal distance between S_1 and S_3.[20] Now, let *m* be the modal distance threshold of an arbitrary safety condition *s*, that is: a world *w* belongs to the class C of worlds in terms of which *s* is cashed out iff the degree of modal distance between *w* and the actual world (in the present case, S_1) is less than *m*. To obtain the expected intuitive results, d_1 should be greater than *m* (so S_2 does not belong to C when S_1 is taken as the actual world, and the belief in *q* is safe in S_1). Now, whichever degree of modal distance d_2 is, it is clear that d_2 is greater than d_1 (astronomically so!), so d_2 should also be greater than *m*. Therefore, a reasonable modal threshold allowing for the belief in *q* to count as knowledge in S_1 would count the belief in *p* and the belief in p_k as safe in S_1 as well.

The previous result is bad enough, since it shows that Safety cannot deal correctly with beliefs in *some* intergalactic propositions, that is, those that involve tickets whose number differs from the winning one in a large number of their digits. But we can show that the problem generalizes to the rest of the intergalactic propositions. This is because the response we temporarily granted with respect to counterfactually different outcomes of isolated subprocesses is not very good either, since we can think of variations of the Intergalactic Lottery Case in which each subprocess has a significantly larger scale than the one specified in the original scenario, and there seems to be no limit in how much that scale can be increased to make the subprocesses large

[20] As indicated in note 4, I am assuming Safety is cashed out in terms of the notion of overall similarity.

enough to have different outcomes only in far-off worlds, according to any reasonable modal distance threshold.

4. Does the Intergalactic Lottery Case Work Against Global Safety?

So far, we have seen that the Intergalactic Lottery Case functions as a counterexample to Safety. Global Safety, however, seems to provide an easy way to block it. According to such principle, a belief an agent has is (globally) safe iff *some* false belief could be formed by the same agent, in the same way, in some C-world. Let $n_@$ be the actual winning ticket's number. Consider now a world *w* in which the number of Annie's ticket is $n_@$ instead of *k*, and everything else in *w* is the same as in the actual world. In *w*, Annie forms a belief in the same way in which she actually formed her belief in the target proposition (that is, that her ticket is not a winner), but the lottery proposition she believes in *w* is different from the target proposition (the former proposition is true in a given world iff $n_@$ is not the winning ticket's number in such world, while the target proposition is true in a given world iff *k* is not the winning ticket's number in that world).[21]

Intuitively, *w* is very close to the actual world, since there are only slight differences needed for Annie's ticket to have a different number (that is, $n_@$) than the number of her actual ticket (that is, *k*). If so, there is a very close world —plausibly, belonging to the relevant class of C-worlds— in which Annie forms a false belief in the same way in which she actually formed her belief in the

[21] The claim that both propositions are different is based on the assumptions a) that the demonstrative term 'my ticket' is a rigid designator and b) that the definite description 'the ticket #k' is not a rigid designator. In my view, (a) is sufficiently supported by the pre-theoretical intuition about the rigidity of demonstratives, while (b) is supported by the intuition that a ticket (understood as the concrete object playing a ticket-role in a lottery) has the number it has contingently. This second assumption might be wrong, if a ticket is an abstract object that is essentially tied to its number. In that case, the scenario proposed is not a possible one.

target proposition that her ticket (that is, ticket #*k*) is not the winner of the lottery. Consider now a specification of Global Safety requiring that there be no C-worlds in which S forms a false belief in the same way in which S actually formed her belief in the target proposition in order for the latter belief to be safe. Given such specification of Global Safety, the target belief is unsafe (thus, being epistemically luckily true), allowing to block the counterexamples proposed so far.

Fortunately, the Intergalactic Lottery Case can be modified to obtain a counterexample to Global Safety. To do so, let us begin by noticing that the process of selecting a winning ticket is, in the scenario of the Intergalactic Lottery Case, a *modally stable* process, that is, a process that could not have easily had a different outcome than the actual one. It is because such a process is modally stable that the target proposition, *p*, could not have easily been false. In contrast, the correlation between the next elements is not modally stable: i) the way in which S's belief in *p* was actually formed and ii) that the belief of S has *p* as its object (it is not modally stable since S could have easily formed a belief in a different proposition in the same way in which S actually formed the belief in *p*). Once we notice this contrast, it is easy to see that all we need to do to obtain a counterexample to Global Safety is to add further details to the description of the way in which S's belief that *p* is actually formed, in such a way that the correlation between (i) and (ii) is modally stable. We can do this in as follows:

> Consider the Intergalactic Lottery case depicted earlier. Suppose additionally that, hundred of thousands of years ago when the Intergalactic Lottery was fist instituted, a flawless technology was created to analyze the genetic code of the citizens of the IFP existing at that time, to obtain (by means of a secret, unchanging and non-random algorithm) a numeral with an astronomical number of digits serving as the lottery ticket of each citizen (thus, each lottery ticket was essentially its number rather than a concrete object). All the

tickets thus generated were significantly different from one another (even if the genetic codes of the ancient citizens did not vary significantly, the algorithm nonetheless generated a radically different ticket for each one). After a citizen died, her ticket was inherited to her direct descendants, if she had any (if she had many, her ticket was inherited to one of them, and the algorithm provided her other descendants a significantly different ticket). This hereditary practice continued for thousands of years up to the present. Given this process for generating tickets, Annie could not have easily had a significantly different ticket. Since the winning ticket could not easily have been different, and since Annie's ticket varied significantly from the winning ticket, Annie could not have easily formed a belief in a false proposition in the same way in which she actually formed her belief that her ticket is a loser.

With these modifications, the Intergalactic Lottery case now functions as a counterexample to both Safety and Global Safety. However, the next revision of Global Safety seems to be sufficient for blocking the present counterexample:

**Global Safety:*

S's belief in *p* is epistemically luckily true if, and only if, S's belief in *p* is **globally unsafe*, that is: there is some C-world in which a belief of an agent S* (possibly different from S) in a proposition *q* is formed in the same way in which S's belief in *p* was formed in the actual world, and *q* is false.

The rationale for using *Global Safety to deal with the previous version of the Intergalactic Lottery Case is captured in the next line of reasoning:

> Any epistemic agent in the Intergalactic Lottery Case could have formed a belief in the lottery proposition about her own ticket in the same way in which Annie formed her belief in the target proposition. This applies to the citizen of the IFP, named Zhee-Vra, who actually has the winning ticket. Had Zhee-Vra formed a belief in the lottery proposition that her ticket is not the winner in the same way Annie's belief in the target proposition was actually formed (all else remaining as in the actual world), she would have ended up forming a false belief. Plausibly, Zhee-Vra could have easily formed her belief in her corresponding lottery proposition in the same way Annie actually formed hers. Thus, some agent (that is, Zhee-Vra!) could have easily formed a false belief in the same way Annie actually formed the target belief. Therefore, the target belief is unsafe according to *Global Safety.

This reasoning makes plausible that *Global Safety can block the version of the Intergalactic Lottery Case presented in this section. But that very same reasoning indicates a way to further modify the scenario in question so that it also works as a counterexample to *Global Safety. For that we only need to suppose that Zhee-Vra simply does not have the ability to make the same reasoning that gave rise to Annie's belief in the target proposition (maybe Zhee-Vra is an infant, lacks the proper conceptual skills to perform the reasoning in question, or lacks the linguistic competence to make self-reference by using a demonstrative term), and that such deficiency is modally stable (perhaps, it is due to some genetic deficiency recurring in all of Zhee-Vra's ancestors). Whatever detail is chosen here to produce the desired version of the Intergalactic Lottery Case, it seems plausible that such details can be worked out into a counterexample to *Global Safety.

Is there any amendment to *Global Safety that avoids the counterexample? Suppose that some such amendment is obtained. Plausibly, it would come from identifying some modal instability in either:

a) the state of affairs making the target proposition true, or

b) the correlation between the way in which the agent's belief in such proposition is formed and other possible beliefs that such agent could have formed in the same way.

However, with sufficient imagination, it always seems possible to make further stipulations to the scenario to increase the modal stability of both (a) and (b), obtaining a counterexample for the amended version of *Global Safety.

Finally, it might be objected that the Intergalactic Lottery Case in its different variations is just too far fetched and designed too intricately to merit serious consideration. But while it involves plenty of detailed science-fictional content, it is conceptually possible, and it does not require something bizarre or outrageous about the cognitive life of the epistemic agents in it. Quite the opposite. From an epistemological point of view, the Intergalactic Lottery is just like any normal lottery, except for the astronomical number of tickets and the astronomical number of their digits. So, despite its fancifulness and intricate design, the Intergalactic Lottery Case shows something important about the conceptual relation (or lack of it) between the epistemic status of lottery propositions and their condition of being safe or unsafe: we can be ignorant of modally-stable lottery propositions.[22]

Moreover, the intuition that an epistemic agent like Annie does not know that her ticket is a loser remains strong despite the detailed and far-fetched fictional content of the scenario. Given

[22] Special thanks to an anonymous referee for a version of this objection.

that Annie's belief is true and (intuitively) justified, the only reason left for why it fails to be knowledge is that it is epistemically luckily true. In that regard, the counterexample might only seem to be a case of safe ignorance. However, it shows the next two claims to be incompatible: i) epistemically luckily true beliefs are unsafe beliefs (that is, Safety or one of its variants is true), and ii) (typical) beliefs in lottery propositions are epistemically luckily true. If beliefs in lottery propositions fail to be knowledge for a reason other than being epistemically luckily true (for example, for being unjustified), the Intergalactic Lottery Case is no counterexample to Safety. If, on the other hand, Safety is false, the Intergalactic Lottery Case is compatible with beliefs in lottery propositions failing to be knowledge because they are epistemically luckily true. The main moral to be drawn from the above cases is that conceptually tying epistemic notions to non-epistemic modal notions, such as safety, is not a safe theoretical move to make.[23]

[23] Thanks to the audiences at McGill University and the participants of the Epistemology seminar at Universidad Nacional Autónoma de México (UNAM) for their feedback. Special thanks to Michael Blome-Tillmann, Melahuac Hernández, Peter Klein, Stephanie Leary, and Lourdes Valdivia for their helpful discussion on earlier drafts.

References

Blome-Tillmann, M. (2020a). Non-Reductive Safety. *Belgrade Philosophical Annual* 33 (33):25-38.

Blome-Tillmann, M. (2020b). Statistical evidence, normalcy, and the gatecrasher paradox. *Mind*, *129*(514), 563–578.

Douven, I. (2002). A new solution to the paradoxes of rational acceptability. British Journal for the Philosophy of Science 53 (3):391-410.

Douven, I., & Williamson, T. (2006). Generalizing the lottery paradox. *The British Journal for the Philosophy of Science*, 57(4), 755–779.

Douven, I. (2012). The lottery paradox and the pragmatics of belief. *Dialectica*, *66*(3), 351–373.

Engel Jr, M. (2010). Epistemic Luck. In Jonathan Dancy, Ernest Sosa & Matthias Steup (eds.), *A companion to epistemology, second edition*. Blackwell. pp. 336-340.

Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford University Press.

Ichikawa, J. J. (2014). Justification is potential knowledge. *Canadian Journal of Philosophy*, *44*(2), 184–206.

Kelly, Kevin T. & Lin, H. (2012). A Geo-logical Solution to the Lottery Paradox, with Applications to Nonmonotonic Logic. *Synthese* 186 (2):531-575.

Kelp, C. (2017). Lotteries and justification. *Synthese*, 194(4), 1233-1244.

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

Pritchard, D. (2005). *Epistemic luck*. Clarendon Press.

Pritchard, D. (2007). Anti-luck Epistemology. *Synthese* 158(3), 277–97.

Pritchard, D. (2008a). Sensitivity, safety, and anti-luck epistemology. *The Oxford Handbook of Skepticism*, 437–455.

Pritchard D. (2008b). Knowledge, Luck and Lotteries. In V. Hendricks (ed.), *New Waves in Epistemology*. Basingstoke: Palgrave-Macmillan.

Pritchard D. (2012). Anti-Luck Virtue Epistemology. *Journal of Philosophy* 109(3), 247–279.

Pritchard D. (2015). Anti-luck Epistemology and the Gettier Problem. *Philosophical Studies* 172(1), 93–111.

Rabinowitz, D. D. (2011). The Safety Condition for Knowledge. *Internet Encyclopedia of Philosophy*.

Ryan, S. (1996). The epistemic virtues of consistency. *Synthese*, *109*(2), 121–141.

Smith, M. (2016). *Between Probability and Certainty: What Justifies Belief Oxford University Press*.

Sosa, E. (1999). How to Defeat Opposition to Moore. *Philosophical Perspectives* 13, 137-149.

Williamson, T. 2000. *Knowledge and Its Limits*. Oxford, England: Oxford University Press.

Yang, D. (2019). What's Wrong with Modal Conceptions of Luck and Risk. *Erkenntnis* 86(4), 773–787.

* Section 2's integration into the dissertation project

There is wide agreement that *safety* is the best candidate to account for epistemic luck. The post-Gettier literature included various accounts of epistemic luck that sooner or later faced serious shortcomings.[24] Yet, safety was an exception among these accounts insofar as it continuously evolved into versions that avoided problems.[25] In particular, safety has garnered support due to having (presumably) successfully evolved to account for our lack of knowledge of lottery propositions: a version of safety requiring (for knowledge) the target belief to be true in "most close worlds" evolved into a version requiring such belief to be true *also* in "all very close" worlds (Pritchard, 2008b).

The perseverance of safety theorists in accounting for our lack of knowledge of lottery propositions reveals their underlying assumption that, while such propositions are highly probable, they cannot but be modally *un*stable. As I have shown in section 2, lottery propositions can be modally stable, and no version of safety can account for why we fail to know them.[26] So, for those lottery propositions that we don't know but are modally stable, are they epistemically luckily true? If the answer is yes, then we need to develop an account of epistemic luck that explains why lottery beliefs are epistemically luckily true and avoids important issues for all previous accounts. As of now, however, there is no such account appearing on the horizon. This constitutes a motivation to examine whether lottery beliefs are justified. I take on this project in sections 3 and 4.

---

[24] See Engel Jr. (2024) for an overview of failed attempts to account for Gettier cases and other cases of epistemic luck.

[25] For example, Hiller & Neta (2007) presented cases of safe ignorance to a local version of safety, which are avoided with a global version of safety.

[26] There is a knowledge-first version of safety, "simple safety" (Blome-Tillmann, 2020), which I do not evaluate in section 2 since it does not fall within the scope of reductive safety theories this section aims to address. According to this version, necessarily, if one knows *p*, one could not easily have been wrong in a similar case. Rather than a being a notion of *overall* similarity, similarity is epistemically restricted to similarity "in those respects that are relevant for knowledge". It is worth exploring whether simple safety can account for why we fail to know the modally stable lottery propositions presented in section 2.

3. Why Are Lottery Propositions Unjustified? : A Critical Survey

1. Introduction

This paper is concerned with the (epistemic) justificatory status of so called ‘lottery propositions’, i.e., propositions reporting that a particular highly probable outcome of a statistical setup obtains. For the present purposes, a lottery proposition is narrowly defined as follows: Given a fair lottery L with n members and one winner, a lottery proposition is a proposition of the form ⟨Ticket #i of L will lose⟩.

Discussion on the epistemic status of lottery propositions arises from a background of puzzles surrounding lotteries. For example, consider H. Kyburg’s (1961) lottery paradox. Take one ticket in a fair one-million ticket lottery with only one winning ticket. Based on the very high odds of such ticket losing, one can come to believe that it will lose. This lottery proposition appears to be justified for us, although without information on the lottery results, we do not seem to know that such ticket will lose. Because the odds are the same for every lottery proposition involved, any of them will be equally justified. Yet, if justification is aggregative (that is if *J*j and *J*ψ, then *J*⟨φ ∧ ψ⟩, where *J* stands for justification), we are then justified in believing that all tickets in the lottery will lose. However, we know –and so, are justified in believing– that not all tickets will lose.

The lottery paradox, and other lottery puzzles, give us motivation to search for a satisfactory account of why we fail to know lottery propositions. If the odds of our ticket losing are so high that it is almost certain that it will lose, why do we fail to know them? After all, we seem to know many propositions about our ordinary world based on evidence that does not make them as extremely probable as lottery propositions - for example, knowing that the convenient store is open based on a neighbor's testimony, or knowing that my colleague's name is Peter based on my memory of having heard his name when we were introduced.

An account of why we do not know lottery propositions identifies a condition necessary for knowledge that lottery propositions do not meet.[27] Two dominant hypotheses of why we do not know lottery propositions are that:

No-Justification

We are not (epistemically) justified in believing lottery propositions based solely on their high odds of being true given the evidence.

Luck

Beliefs in lottery propositions happen to be true by sheer epistemic luck, (i.e., luck involved in arriving at a true belief that prevents it from being knowledge, even if it is justified).

This paper focuses on the first hypothesis, and evaluates accounts based on it to explain why we fail to know lottery propositions. Because the only type of justification we consider is inferential, 'justified' is restricted to inferential justification throughout this paper.

Given that the sole justification of lottery propositions is inferentially based on purely statistical evidence, if we assume that a proposition is inferentially based on the evidence just consists in it being highly likely given the evidence, lottery propositions seem to be justified, since they can be made as likely as we want depending on the size of the lottery (even more likely than

[27] The assumption that justification is necessary for knowledge is maintained throughout this paper. It is important to notice that accepting justification as necessary for knowledge does not commit us to accepting that knowledge analyzes into justified true belief plus an additional condition that converts said belief into knowledge. There are knowledge-first views according to which justification is a necessary condition for knowledge even though justification is understood in terms of knowledge (for example C. Kelp's view (2016, 2017a), here discussed in section IV).

other ordinary propositions that we intuitively consider to be justified). However, there are strong reasons (considered in the next section) for thinking that lottery propositions are not justified, thus motivating our focus on No-Justification. [28]

This paper critically surveys some of the main views proposed in the literature on the justificatory status of lottery propositions that explain, or suggest an explanation of, No-Justification. Views proposing a condition for inferential justification that is reducible to logical and probabilistic notions -henceforth called "probabilistic" views- are surveyed in Section III. This type of view is suggested by the fact mentioned earlier, that the support of lottery propositions is solely inferentially based on statistical evidence, so relying on purely logical and probabilistic notions could be thought to suffice in explaining No-Justification. First, two views representative of this kind are surveyed: Sharon Ryan's (1996) view, according to which lottery propositions are not justified because they belong to a set with a member that is false but which we are unable to identify, and Igor Douven's (2002) view, according to which lottery propositions are not justified given that they belong to a self-undermining set. Then, a case against them proposed by Douven and Williamson (2016) is examined and evaluated. It is argued that although such case succeeds against Ryan and Douven's proposals, it does not succeed against other probabilistic views. Since I have argued elsewhere (Ochoa, 2021) against Douven and Williamson's general proof attempting to show that probabilistic views are untenable, I conclude that probabilistic views might still be a feasible option. Non-probabilistic views are surveyed in Section IV. Douven's (2010) pragmatic account, Ichikawa's Knowledge-First account (2014), Kelp's virtue epistemology account (2017a,b), and Smith's (2010,2016) normic support accounts will be examined. I conclude that,

[28] An alternative way to motivate No-Justification is that accepting it provides a way to solve the lottery paradox. Yet, some authors welcome the idea that we are justified in believing lottery propositions (e.g., Foley, 1987, and McGlynn, 2012). To solve the lottery paradox, this position would most likely need to reject that we are not justified in believing the conjunction of all lottery propositions of a given lottery, by denying the assumption that justification is aggregative.

among the non-probabilistic views discussed, Smith's view best serves the purpose of explaining why we do not know lottery propositions on the basis of the hypothesis that we lack justification for them.

2. Some Preliminaries

Recently, some independent arguments in favor of No-Justification have been provided, such as the following ones:

> A1. Consider an arbitrary lottery proposition ⟨*P*⟩. No person is justified in believing a proposition of the form 'ϕ and I don't know that ϕ' (since they are Moore paradoxical). So, I am not justified in believing ⟨*P* and I don't know that *P*⟩. But then, either I am not justified in believing ⟨*P*⟩ or I am not justified in believing ⟨I don't know that *P*⟩ (if I were justified in believing both propositions, I would be justified in believing their conjunction). However, I know ⟨I don't know that *P*⟩ (since I know that lottery propositions are unknown), so I am justified in believing it (since knowledge entails justification). Therefore, I am not justified in believing *P*. Since P is an arbitrary lottery proposition, it follows that I am not justified in believing lottery propositions. Kelp (2017)

> A2. Justified beliefs falling short of knowledge are at least candidates for knowledge, that is: if different non-justificatory (e.g., non-evidential) conditions were met, a justified belief would be knowledge (for example, a perceptual belief justified on the basis of sensory experience in uncommon unreliable circumstances, would be knowledge in ordinary reliable circumstances). However, beliefs in lottery propositions are not candidates for knowledge, since they would only become knowledge if a better justification were

available (e.g., reliable testimony that my ticket is the winner). Therefore, beliefs in lottery propositions are not justified. (Smith, 2021: 98-100)

A3. For an agent to have a justified belief in a given proposition, it *should* be surprising if it turns out to be false, since justification is meant to help avoid false beliefs. However, it should not be surprising that a lottery proposition turns out to be false (it might be surprising that *my* ticket won, but it shouldn't be, since it is like any other lottery ticket and it is not reasonable to be surprised that a certain ticket won). (Smith, 2021: 100-102)

A4. Paradigmatic cases of justified beliefs (e.g., usual perceptual beliefs) are not risky in two senses, that is: every circumstance in which they are false and formed in the same way as in the actual world is a circumstance that is i) far away from the actual world (so, the beliefs are safe), and ii) highly abnormal (e.g., they need to involve vivid hallucinations). In contrast, beliefs in lottery propositions could easily be false in normal circumstances. So, most plausibly, beliefs in lottery propositions are not justified. (Smith, 2021: 102-106)

A5. Justified beliefs are epistemically resilient against contrary testimonial evidence. For example, the evidence in favor of justified perceptual beliefs is only slightly weakened (insufficient to make it reasonable to stop holding them) when presented with testimony to the contrary. However, beliefs in lottery propositions are not epistemically resilient against contrary testimony (overhearing a passerby telling his friend that the lottery results are out and that ticket #543 won is sufficient to suspend judgement). Therefore, most plausibly, lottery propositions are not justified. (Smith, 2021: 106-109)

If we accept these arguments, then the question that arises is that of *why* lottery propositions fail to be justified (that is, why is No-Justification true?). Answering this question is crucial for the plausibility of the previous arguments, since not being able to find such an explanation despite surveying a large sample of possible answers would raise doubts on whether the arguments are really cogent. Thus, a full defense of No-Justification is only complete after such an explanation is found.

Some accounts examined in the following sections have the explicit goal of explaining No-Justification (Ryan 1996, Smith 2010, and Kelp 2017). However, since there are few proposals of this nature, it is better to extend the scope of our survey to accounts not having that explicit goal. A natural way of making such an extension results from a brief examination of the next argument against No-Justification:

> Lottery propositions are extremely probable given the evidence (certainly, they are much more likely than most propositions that we ordinarily accept as justified). Since having a high probability given our evidence is sufficient for being well-supported by it, lottery propositions are well-supported by our evidence. Thus, we are (inferentially) justified in accepting them.

This argument relies on the next premise:

> Sufficiency
>
> Having a high probability (even if it is below 1) is sufficient for being well-supported by the available evidence

Although Sufficiency is quite plausible at an intuitive level, no apparent reason forces us to accept it in all its generality, so restricted versions of it that are compatible with No-Justification might be true. More importantly, there is a suggestive relation between a suitable restriction to Sufficiency and an adequate explanation of No-Justification, which I turn to next.

Firstly, any explanation of No-Justification provides a motivated way to restrict Sufficiency. At a minimum, such an explanation should entail that there is a property satisfying the following schema:

NEC: N is a property such that no lottery proposition instantiates N and only propositions instantiating N are justified.

Secondly, any suitable restriction to Sufficiency entails that there is a property satisfying the following schema:

SUF: S is a property such that every lottery proposition instantiates S, and every highly likely proposition that does not instantiate S is justified.

Notice now that if a given property satisfies NEC, then lottery propositions are counterexamples to Sufficiency, since they are highly probable but fail to be justified for lacking such property. For this very reason, whatever condition is added to having a high probability so as to obtain a suitable restriction of Sufficiency must be incompatible with instantiating a property satisfying NEC (otherwise, the added condition would not avoid being a counterexample to the obtained restriction, in the same way in which lottery propositions are counterexamples to Sufficiency). Thus, a way to motivate a suitable restriction to Sufficiency is to define it in terms of a property satisfying SUF such that its negation satisfies NEC, guaranteeing that no highly likely proposition instantiating the property is a counterexample to the proposed restriction.

Given the previous relation between explanations of No-Justification and suitable restrictions to Sufficiency, this survey also considers various proposals restricting Sufficiency to assess whether they can be used to provide a property satisfying NEC.

Before beginning the survey, let me briefly describe two explanations of No-Justification this paper does not address (due to space constraints, they are left for future discussion). The first one consists in the view that, however high the probability of a lottery proposition given the evidence, it is *never* high enough for it to be justified. Unfortunately, taken as an explanation of No-Justification, this restriction has very strong consequences. Notice that the probability of lottery propositions can be as close to 1 as we want, provided that the relevant lottery is large enough. So, for this response to work as generally as needed, it would need to entail the claim that any proposition with a probability below 1 given the evidence is not justified. This claim is called 'Infallibilism'. Although Infallibilism has the advantage of being compatible with a very simple understanding of inferential justification in terms of high probability, it is extremely costly. This view has the skeptical implication that inductive reasoning is never strong enough to produce knowledge, being at odds with the widely held view that most of our empirical knowledge is inductive.

The second one begins by distinguishing between two senses of 'belief', a weak one and a strong one (also called 'outright belief'), depending on the strength of the normative standards for rationally holding each type of belief-state. Based on such a distinction, it is held that lottery propositions meet the normative standards for weak belief, but not the standards for strong belief, while assuming that the latter are the same as those governing knowledge. We can call this view 'Compatibilism', since it yields a compatibilist result between seemingly incompatible intuitions about whether beliefs in lottery propositions are justified, by entailing that there is a sense of 'belief' in which believing lottery propositions is justified, and a sense in which it is not. Yet, Compatibilism thus understood counts as an explanation of No-Justification, since it identifies a sense of ‘belief’ (outright belief) that is greatly significant for epistemology given its close relation to knowledge, and under this sense beliefs in lottery propositions are not justified.

3. Probabilistic Accounts of Why Lottery Propositions Lack Justification

In this section probabilistic accounts will be examined. Let us remember that probabilistic accounts are accounts that ultimately explain the lack of justification of lottery propositions exclusively in terms of logical and probabilistic notions. Throughout this section I will be referring to a class of lottery propositions defined as follows: LOT = { x | x = ⟨Ticket #i of L will lose⟩, for every i such that $1 \leq i \leq n$}, where L is a fair lottery with only one winning ticket and all tickets having equal odds of losing. LOT is intentionally underspecified, since the following discussion is intended to apply to any lottery with the same features.

1. Douven: Belonging to a Self-Undermining Set

An explanation of No-Justification can be extracted from Igor Douven’s (2002) restriction to Sufficiency (i.e., having a high probability is sufficient for being well-supported by the available evidence), saying that a proposition is justified if, in addition to having a high probability, it does not belong to a “self-undermining” set. A set Φ is self-undermining iff, for each member x of Φ, the probability of x given E is above t, but the probability of x given the union of E with Φ-x is not above t.[29] For example, consider a fair lottery of 1000 tickets with only one winner, so LOT = $\{p_1, p_2, p_3\ldots, p_{1000}\}$, and suppose that the threshold for being supported by E is having a probability

[29] See Douven and Williamson (2006:759). For a more specific version, see Douven (2002:396).

given E above 0.5. The probability of p_1 given E is 0.999, but its probability given E∪{p_2,p_3,…,p_{999}} is 0.5, and its probability given E∪{p_2,p_3,…,p_{1000}} is 0. The same applies to any other member of LOT, so LOT is self-undermining, and no member of LOT meets Douven's restriction. Thus, Douven's restriction of Sufficiency avoids concluding that lottery propositions are justified.

Can a property that is necessary for justification be obtained from Douven's restriction to Sufficiency? Arguably, it might. To do so, let us first notice that the support of lottery propositions by the evidence is non-deductive, so it is *defeasible*. The support of a proposition P by a set of evidence E is defeasible when a non-deductive argument having E's members as premises and P as conclusion strongly supports P, but it stops doing so when certain premises are added to it. Consider, for example, the next pair of arguments:

A1

1. Peter is a philosopher
2. Most philosophers are boring

$\therefore_{prob}$ C. Peter is boring

A2

1. Peter is a philosopher
2. Most philosophers are boring
3. Peter plays in a band
4. Almost no philosopher playing in a band is boring

$\therefore_{prob}$ C. Peter is boring

Intuitively, argument A1 strongly supports its conclusion, while A2, which differs from the first one in containing more premises, does not support it. Thus, 1 and 2 provide defeasible support to C since by adding 3 and 4 they stop supporting it. If E = {1, 2, 3, 4}, we might think that E supports C if we only considered 1 and 2, since its subset {1, 2} supports C, but when we consider all the members of E it is clear that E does not support P, since {1, 3, 4} actually supports

¬C. Here, a subset of the available evidence supports a proposition, but a different subset with the rest of the available evidence provides counterevidence undermining such proposition's support when combined with the first subset. The defeasible character of non-deductive support by E makes plausible the next general principle about justification:

Defeasibility Principle

If P is justified, then P is supported by E, and it is not the case that a subset of E supports P but its union with the rest of the members of E does not support P

Defeasibility Principle prevents proposition C from being justified when supported by evidence E = {1, 2, 3, 4} in argument A2. In this case, only part of the evidence supports C, i.e., {1, 2}. Yet, {1, 2}∪{3, 4} does not. Notice, however, that lottery propositions in LOT satisfy Defeasibility Principle. Consider the way in which they are supported by E, expressed in the next statistical syllogism (where L is a fair lottery with 1000 tickets and only one winner):

L1. 99.9% of tickets belonging to L are losing tickets

L2. Ticket #*i* belongs to L

Therefore (most likely),

L3. Ticket #*i* is a losing ticket

In this case, the available evidence E consists in our known information L1 and L2. Notice that *i* can be any lottery ticket number in L, so any member of LOT is supported by E via an argument of this form. For Defeasibility Principle to prevent L3 from being justified, there would need to be subsets of E that could be extracted from L2 and L3, such that their union prevented L3 from being supported. There are no such subsets. Thus, Defeasibility Principle does not prevent L3 from being justified.

Of course, this is not to say that lottery propositions are not defeas*ible*. Even when the above argument strongly supports L3, new information added as premises would defeat L3's support by {L1, L2}. For example, suppose that, according to the latest news, ticket #*i* won the lottery. This piece of information is not part of the evidence, so L3's support is not actually defeated, but it is nonetheless defeasible.

Although Defeasibility Principle does not prevent lottery propositions from being justified, lottery propositions are importantly similar to propositions rendered unjustified by such principle. With respect to the latter, there is information *contained in the evidence* that renders them unsupported. Similarly, with respect to lottery propositions, there is information that is *supported by the evidence* and which, along with the evidence, renders them unsupported. Such information consists in other propositions in LOT. For example, suppose E supports a lottery proposition p_1. Just as it supports p_1, it equally supports p_2, p_3,…, and p_{1000}. But {p_2,…,p_{999}} or {p_2,…,p_{1000}}, together with E, make the probability of p_1 equal to or below the 0.5 threshold of support. The above similarity motivates the following variant of Defeasibility Principle to explain why lottery propositions fail to be justified:

Defeasibility Principle*:

If P is justified, then P is supported by E, and it is not the case that E supports each member of a set of propositions ψ such that E ∪ ψ does not support P

Lottery propositions fail to satisfy the consequent of this principle. Consider again how such propositions belong to a self-undermining set, according to Douven. For each lottery proposition p_i in LOT, p_i is supported above t by E but the union of E with the rest of the propositions supported by E undermines p_i's support to a degree that is not above t. So, by belonging to a self-undermining set, p_i is supported by E, but there is always a set ψ-P -i.e., the relevant subset of LOT- whose members are each supported by E and whose union with E does not support p_i. For example, when P=p_1 and ψ−P={p_2,…,p_{1000}}, E supports each of ψ's members, p_2,…,p_{1000}, above t, and E∪{p_2,…,p_{1000}} does not support p_1 above t. Therefore, the second conjunct of Defeasibility Principle*'s consequent is not satisfied and, as a result, such consequent is false for lottery propositions. This allows us to conclude that they fail to be justified. Thus, if Defeasibility Principle* is a plausible principle of justification, it allows us to explain why lottery propositions fail to be justified in virtue of belonging to LOT.

2. Ryan: Belonging to a Set with a Member that Is False Without Knowing Which Member It Is

We can find an explicit attempt to explain No-Justification in Sharon Ryan (1996). Her diagnosis is that there is a form of justified counterevidence to each lottery proposition preventing us from being justified in believing them.

Ryan's explanation arises from analyzing a particular situation in which lottery propositions are: despite having good reason to believe each member of LOT, we also know that one of its members is false but we do not know which one it is. When this known fact is added to our total evidence, it bears on the justificatory status of every member of LOT by calling each of them into question without the means to decide between them (130-1). To motivate her account, Ryan describes two cases analogous to the lottery case:

Book mystery case: I share office with Mark and Ned, about both of whom I have equally good reason to think they are reliable and have not stolen a book missing from my desk. But later I come to know that they were the only people present in the office during the time the book disappeared. I know one of them took it, but do not know who it was. Intuitively, this piece of information is counterevidence to my beliefs that Mark did not take the book and that Ned did not take the book, calling each one of them into question.

Furniture-throwing case: Smith and Jones have very good reason to think that their furniture will remain in their room after they leave it today, that it will remain in the room after they leave tomorrow, that it will remain in the room after they leave the day after tomorrow, ... and so on. But suppose they discover, after years of therapy, that Jones has a syndrome that will make him frantically throw away all the furniture out of the window within the next one million days. So, they know that on one of the next million days the furniture will not remain in the room. Intuitively, this new piece of information is counterevidence to their beliefs, each about one of the next million days, that the furniture will remain in the room. It calls each of such beliefs into question.

In the above cases, the agents have good reason to believe each member of the following sets of propositions:

B: {⟨Mark did not take the book⟩, ⟨Ned did not take the book⟩}

F: {⟨the furniture will stay in place on day 1⟩, ⟨the furniture will stay in place on day 2⟩,…, ⟨the furniture will stay in place on day 1,000,000⟩}

Nonetheless, they also know that at least one of the members of B and of F is false. They know the propositions c_B= ⟨either Mark or Ned took the book⟩ and c_F= ⟨one day of the 1,000,000 days the furniture will not stay in place⟩.

Intuitively, despite having good reason to believe each member of B and F, knowing that c_B and c_F prevents an agent from being justified in believing them. Despite having good reason to believe that Mark did not take my book, knowing that either Mark or Ned took it makes me no longer justified in believing that he didn't take it. Similarly, despite having good reason to think that the furniture will remain in the room tomorrow, learning that Jones will throw it away one of the following million days makes Smith not be justified in believing it. These intuitions are codified in the next principle:

> Avoid-Falsity Principle
>
> For any set Ψ of propositions, if (i) S has good reason to believe each of Ψ's members and (ii) S has good reason to believe that at least one member of Ψ is false (or to suspend judgement about at least one of them), then S is not justified in believing any member of Ψ.

Ryan defines "having good reason to believe p" as one's evidence supporting p more than it supports ¬p. Just as in the book and furniture cases Ψ=B and Ψ=F, in lottery cases Ψ=LOT. As part of our evidence, we have equally strong reasons for each member of LOT. But once the known proposition c_{LOT}= ⟨one of the 1000 tickets is not a loser⟩ becomes part of our total evidence, such evidence no longer justifies any member of LOT. Thus, Avoid-Falsity delivers the result that lottery propositions are not justified.

3. **Douven and Williamson's Case against Some Probabilistic Accounts**

Defeasibility Principle* and Avoid-Falsity Principle aim to provide necessary conditions for justification that lottery propositions do not meet. As explained in section 2, if there is a property necessary for justification that lottery propositions lack, the condition of possessing it can restrict Sufficiency. A restricted version of Sufficiency will then prevent us from concluding that lottery propositions are justified. Such a version will instantiate the following schema, where φ is any proposition and t the threshold for justification:

Restricted Sufficiency

If $t < Pr(\varphi) < 1$, and it is not the case that φ has a defeater *D*, then φ is justified[30]

Following Douven and Williamson (2006), a defeater is any property such that possessing it overrides or undermines φ's justification. For example, consider Douven's (2002) proposed property of belonging to a self-undermining set. If *D* is the property of belonging to a self-undermining set, lottery propositions possess *D*. But having such property renders lottery propositions unjustified, as seen in section A. Similarly, lottery propositions possess *D* if *D* is Ryan's (1996) property of belonging to a set such that the probability of one of their members being false given the evidence is above t (assuming $t < 0.5$). But lottery propositions are not justified because the evidence supports above 0.5 that one of the members of LOT is false. Such

[30] Douven and Williamson (2006) formulate the above schema using the notion of rational acceptability rather than in terms of the notion of (inferential) justification. The authors do not provide a definition of 'rational acceptability', but rather "rely on the reader's informal understanding of the notion"(757). Pretheoretically, our beliefs are rationally acceptable when they are supported by reasons good enough to authorize us to accept or assent to them. The notion of justification can informally be understood in the same way. Thus, I will use both notions interchangeably.

instantiations of *D* yield two versions of Restricted Sufficiency, where lottery propositions do not meet the antecedent of such principle.

In "Generalizing the Lottery Paradox" (2006), Douven and Williamson (DW, to abbreviate) provide an argument showing that for Douven's (2002) and Ryan's accounts of a defeater, Restricted Sufficiency fails. They argue that when *D* is identified with Douven and Ryan's definitions, *any* non-lottery proposition whose probability is above *t* and below 1 will have *D*.[31] Consequently, Restricted Sufficiency will not allow us to infer, for any non-lottery proposition (justified or not), that it is justified.

DW's argument is as follows. Let us call the properties that Douven and Ryan take to be defeaters 'D_D' and 'D_R', respectively. Consider an arbitrary proposition ψ such that $t < \Pr(\psi) < 1$ and, for a lottery *l* with a large enough number of tickets *n* such that $1-(1/n) > t$. Now, consider the set $L = \{ \phi_i \mid 1 \leq i \leq n \wedge \phi_i = \langle \text{Ticket } \#i \text{ of } l \text{ loses} \rangle \}$ (that is, the set of all singular lottery propositions about each ticket of lottery *l*). Since exactly one ticket of *l* wins, for any *i*:

i) $\Pr(\phi_i) = 1-(1/n) > t$.

And, since a disjunction is at least as probable as its disjuncts, we have that, for any *i*:

ii) $\Pr(\langle \neg\psi \vee \phi_i \rangle) > t$.

Given this, it can be shown that ψ has D_D and D_R. Consider now the set $\Gamma = \{\chi \mid 1 \leq i \leq n \wedge \chi = \langle \neg\psi \vee \phi_i \rangle\}$ (the set of propositions of the form $\langle \neg\psi \vee \phi_i \rangle$). Notice that ψ is a member of $\{\psi\} \cup \Gamma$, and all its members have a probability above *t*. We now have that:

31 DW present the same problem for their reconstruction of Pollock's account of a defeater: the property of being a member of a minimally inconsistent set of propositions, where each proposition has a probability above *t* (2006: 759-760).

ψ has D_D: Since one ticket is guaranteed to win, it is part of the background knowledge that some ϕ_i is false, so $\Pr(\psi| (\neg\psi \vee \phi_1), (\neg\psi \vee \phi_2),..., (\neg\psi \vee \phi_n)) = 0$.[32] But $\Pr(\psi| (\neg\psi \vee \phi_1), (\neg\psi \vee \phi_2),..., (\neg\psi \vee \phi_n)) = \Pr(\psi|\Gamma)$. Therefore, $\Pr(\psi|\Gamma) = 0 < t$. Also, for every *i*, $\Pr((\neg\psi \vee \phi_i)| (\neg\psi \vee \phi_1),..., (\neg\psi \vee \phi_{i-1}), (\neg\psi \vee \phi_{i+1}),..., (\neg\psi \vee \phi_n), \psi) = 0$.[33]. Therefore, $\{\psi\}\cup\Gamma$ is a self-undermining set and so ψ has D_D.

ψ has D_R: Since $\{\psi\}\cup\Gamma$ is inconsistent, the probability that all its members are true is 0, so it is below *t*, yet the probability of each one of them is above *t*. Therefore, ψ has D_R.

The above case shows that the identification of *D* with D_D or D_R is unacceptable, since there surely are non-lottery propositions that are justified because they are highly probable and have no justification-defeating property. Given Restricted Sufficiency, we will only be able to conclude of propositions with a probability of 1 that they are justified.

Notice that the failure of Restricted Sufficiency with D_D or D_R does not amount to the claim that this principle, together with D_D or D_R, makes any proposition with probability below 1 and above *t* fail to be justified. For ψ to fail to be justified, ψ not having D_D or not having D_R need to be necessary conditions for justification, not only restrictions to Sufficiency. DW's case only prevents us from inferring that ψ is justified.[34]

---

[32] This is proved as follows:

1. $((\neg\psi \vee \phi_1) \wedge (\neg\psi \vee \phi_2),..., \wedge (\neg\psi \vee \phi_n)) \Leftrightarrow (\neg\psi \vee (\phi_1 \wedge \phi_2,..., \wedge \phi_n))$ (where '$\Leftrightarrow$' stands for logical equivalence)

$\therefore$ 2. $\Pr(\psi| (\neg\psi \vee \phi_1) \wedge (\neg\psi \vee \phi_2),..., \wedge (\neg\psi \vee \phi_n)) = \Pr(\psi| \neg\psi \vee (\phi_1 \wedge \phi_2,..., \wedge \phi_n))$ (from 1).

3. $\Pr(\phi_1 \wedge \phi_2,..., \wedge \phi_n) = 0$ (since one ticket is guaranteed to win)

$\therefore$ 4. $\Pr(\psi| \neg\psi \vee (\phi_1 \wedge \phi_2,..., \wedge \phi_n)) = \Pr(\psi| \neg\psi) = 0$ (from 3).

$\therefore$ 5. $\Pr(\psi| (\neg\psi \vee \phi_1), (\neg\psi \vee \phi_2),..., (\neg\psi \vee \phi_n)) = 0$ (from 2 and 4).

[33] The proof is as follows:

1. $((\neg\psi \vee \phi_1) \wedge,..., \wedge (\neg\psi \vee \phi_{i-1}) \wedge (\neg\psi \vee \phi_{i+1}),..., \wedge (\neg\psi \vee \phi_n) \wedge \psi) \Leftrightarrow ((\neg\psi \vee (\phi_1 \wedge,..., \phi_{i-1} \wedge \phi_{i+1},..., \wedge \phi_n)) \wedge \psi)$.

2. $\Pr((\neg\psi \vee \phi_i)| ((\neg\psi \vee (\phi_1 \wedge,..., \phi_{i-1} \wedge \phi_{i+1},..., \wedge \phi_n)) \wedge \psi) = 0$.

$\therefore$ 3. $\Pr((\neg\psi \vee \phi_i)| (\neg\psi \vee \phi_1) \wedge,..., \wedge (\neg\psi \vee \phi_{i-1}) \wedge (\neg\psi \vee \phi_{i+1}),..., \wedge (\neg\psi \vee \phi_n) \wedge \psi) = 0$.

[34] DW not only consider not having D_D or not having D_R as sufficient conditions for justification, but also as necessary conditions. Otherwise, they would not be able to conclude from their case that "not only lottery propositions, but all propositions having non-perfect probability fail to qualify as rationally acceptable [i.e., justified]", given Ryan and Douven's proposals of a defeater (760).

Nonetheless, as explained in section 2, a good way to motivate a restriction to Sufficiency is by defining such restriction in terms of a property that lottery propositions lack and only propositions instantiating it are justified. Not having D_D and not having D_R seemed to be acceptable candidates for such property, insofar as they explain No-Justification: lottery propositions are not justified because they do not meet the necessary condition of not having D_D or not having D_R (see 3.1 and 3.2). In any case, the upshot in light of DW's proof is that if we wish to maintain these conditions as necessary for justification, we cannot use them to restrict Sufficiency, since any arbitrary proposition with probability above *t* and below 1 will meet them. This is a bad enough and undesirable result.

4. **Reconsidering Probabilistic Accounts**

While DW's case generates a problem for Douven's and Ryan's accounts, it fails to do so for similar, plausible accounts. Consider the following principle, inspired in Ryan (1996):

> Equal Likelihood Principle
>
> For any set Ψ of propositions, if (i) each of Ψ's members has a probability above *t* and below 1 given S's total evidence E, (ii) all members of Ψ are *equally* likely given E, and (iii) the probability that one member of Ψ is false given E is above *t*, then S is not justified in believing any member of Ψ.

This principle can account for why Non-Justification is true. Consider again LOT= { x | x = ⟨Ticket #i of L will lose⟩, for every i such that $1 \leq i \leq n$}. All members of LOT are individually highly likely given the evidence, so they meet (i). Since the members of LOT are propositions about tickets in a fair lottery, they are all equally likely given the evidence; that is, they meet (ii). Finally, our evidence renders it above *t* likely that one of the members of LOT is false; that is, (iii) is met. This follows directly from our knowledge that one of them is false.

Just as Avoid-Falsity, Equal Likelihood can explain why in Ryan's book and furniture cases the subjects fail to be justified in believing the target propositions, which can be any member of the following sets:

B: {⟨Mark did not take the book⟩, ⟨Ned did not take the book⟩}

F: {⟨the furniture will stay in place on day 1⟩, ⟨the furniture will stay in place on day 2⟩,…, ⟨the furniture will stay in place on day 1,000,000⟩}

Just as all members of LOT are equally likely, so are all members of B and of F. With respect to B, we know that either Mark or Ned took the book, and no one else. Since the evidence supports each hypothesis equally, both are equiprobable given the evidence. Similarly, we know that the furniture will leave the room one of the following million days, and there is no reason to suspect it will occur on a specific day over the others. Therefore, given our evidence, the likelihood of the furniture staying in place on any of those days is the same. Therefore, Ryan's book and furniture cases meet condition (ii) of Equal Likelihood. Additionally, they meet conditions (i) and (iii), for these are precisely the two sufficient conditions for justification failure in Avoid-Falsity.

Furthermore, Equal Likelihood can be motivated in the same way as Avoid-Falsity and with these same cases. As seen in section B, it is intuitively the case that the subjects in both the book and furniture scenarios are not justified in believing the target propositions. These scenarios are analogous to the lottery case with respect to (i)-(iii) in Equal Likelihood. So, it is reasonable to think that lottery propositions are not justified because they meet conditions (i)-(iii). In other words, Equal Likelihood is a principle that provides a plausible account of No-Justification – at least as plausible as Avoid-Falsity.

Now, let us see why DW's case poses no problem for an account of No-Justification based on Equal Likelihood. From Equal Likelihood we can extract an overriding property D_E of lottery propositions in LOT, preventing them from being justified: they belong to a set such that a) all its members are equally likely given the evidence and b) the probability that one of the members is false given the evidence is above t (assuming $t < 0.5$). DW's proposed case (in section C) will generate a problem for an account based on Equal Likelihood only if their arbitrary proposition ψ has property D_E. Remember DW's proposed set, $\{\psi\}\cup\Gamma$, where Γ is the set of all propositions of the form $\langle\neg\psi \vee \phi_i\rangle$ and ϕ_i stands for any lottery proposition belonging to LOT. Condition (b) is met, since $\{\psi\}\cup\Gamma$ is inconsistent because at least one of its members is false. But condition (a) is not met, since assuming equal likelihood of all members of $\{\psi\}\cup\Gamma$ has unwelcome results, as shown below.

Let us begin by considering that for a proposition P to be justified based on evidence E, E should support P more than it supports ¬P. If we take probability as a measure of the degree of support, the probability of P should be higher than the probability of ¬P, that is, above 0.5. But even if we suppose that the probability of P is 0.5, a problem arises for the assumption that ψ and the rest of the members of $\{\psi\}\cup\Gamma$ are equiprobable. Consider the following reductio:

1. $Pr(\psi) = ½$	*Ass.*
2. $Pr(\neg\psi \vee \phi_i) = Pr(\psi)$	*Ass. Equiprobability for reductio*
3. $Pr(\neg\psi \vee \phi_i) = Pr(\neg\psi)+Pr(\phi_i) - Pr(\neg\psi)*Pr(\phi_i)$	*Def. by general addition rule, multiplication rule (probability calculus)*
4. $Pr(\psi) = a$	*Arbitrary constant*
5. $Pr(\phi_i) = b$	*Arbitrary constant*

6. $Pr(\neg\psi) = 1\text{-}Pr(\psi)$ *Def. by negation rule (probability calculus)*

7. $Pr(\neg\psi \vee \phi_i) = (1\text{-}a)+b - (1\text{-}a)*b$ *Substitution (3-6)*

8. $Pr(\neg\psi \vee \phi_i) = a$ *Transitivity of = (2,4)*

9. $a = (1\text{-}a)+b - (1\text{-}a)*b$ *Transitivity of = (7,8)*

10. $2a+ab = 1$ *Arithmetic equivalence (9)*

11. $2(1/2)+1/2b = 1$ *Transitivity of = (1,4) and substitution in 10*

12. $1+1/2b = 1$ *Arithmetic equivalence (11)*

13. $b = (1\text{-}1)/\text{-}1/2 = 0$ *Arithmetic equivalence (12)*

14. $\neg(b=0)$ *Ass. about lottery propositions*

15. $\bot$ *(13,14)*

These derivations show that even if the threshold for justification is lowered to a bare minimum of 0.5 probability, ψ cannot be equiprobable with ($\neg\psi \vee \phi_i$) while meeting such threshold, on pain of rendering all lottery propositions (i.e. ϕ_i) zero percent likely. Of course, our background knowledge that exactly one ticket will win in a large lottery excludes this possibility (in fact, it is expected that lottery propositions are highly likely as well).

The above result establishes a difference between the members of LOT and the members of $\{\psi\}\cup\Gamma$. All lottery propositions in LOT are equiprobable, whereas not all members of $\{\psi\}\cup\Gamma$ are equiprobable. Thus, ψ does not possess overriding property D_E. DW's case does not present a problem for an account of No-Justification based on Equal Likelihood.[35]

[35] It could be argued that condition (ii) of the antecedent of Equal Likelihood (that is, that all members of Ψ are *equally* likely given E) is not essential to the lottery paradox, as this paradox also arises at least in cases in which the likelihood of the members of Ψ slightly varies (e.g., where some members are be .999 probable while others .8, it can still be the case that the probability of their conjunction is less than *t* and hence, their conjunction will be unjustified. In such cases, lottery propositions do not possess D_E (extracted from Equal Likelihood) insofar as they are not equiprobable, but they could (arguably) still be rendered unjustified (as Ryan does, for example). Thus, D_E is too strong to use in a solution to the paradox. If so, then the fact that DW's criticism is unsuccessful with respect to a solution to the paradox appealing to D_E may not diminish the force of such criticism. Now, I grant that D_E is too strong to use in a general

Finally, DW present a general argument based on their case against Ryan and Douven's accounts. DW generalize to any defeating property *D* of propositions defined in logical and probabilistic notions, which they dub "a structural property". [36] Their aim is to show that if there is such a justification-defeating property, it cannot be structural, as D_D and D_R are. They assume that *D* is structural and, by reductio, conclude that no property of this kind can be used to solve the lottery paradox by restricting Sufficiency. In "The Aggregativity of Justification and De-Generalizing the Lottery Paradox"(2021), I argue that their general proof is better understood as a reductio of one of the epistemic assumptions of the proof: that justification is "aggregative" –i.e., that the justification of a pair of propositions is sufficient for the justification of their conjunction–, rather than as a reductio of an overriding property being structural.[37] If this is correct, there is good reason to suspend judgement on the epistemic implication of their proof, that is, that any account of justification based on purely probabilistic notions is incorrect.

The moral of DW's case, along with their general proof, is to abandon purely probabilistic accounts of justification. Since (on my view) their general proof fails to problematize such accounts and, as shown above, there are purely probabilistic ones that DW's case does not

---

solution to the paradox. Note, nonetheless, that paradigmatic cases in which the paradox arises are precisely those in which all members of the relevant set of lottery propositions are equally likely and the probability that one of the members is false is above *t*, that is, precisely those in which lottery propositions have D_E (regardless of whether D_E is in fact an overrider). Thus, while DW's criticism is successful with respect to some cases in which the paradox arises, it is limited insofar as it is unsuccessful with respect to other, paradigmatic cases of the paradox. Accepting DW's criticism would amount to accepting that some solutions to the paradox fail (such as Ryan's), *as long as* those solutions do not apply to cases in which the relevant set of lottery propositions are equiprobable. I thank Ulf Hlobil for discussion on the relevance of Equal Likelihood in solving the lottery paradox.

[36] See DW(2006:762) for the technical definition of "structural property".

[37] DW provide a proof to generate a dilemma: if *D* is structural, either Restricted Sufficiency amounts to the trivial claim that having a probability of 1 is sufficient to be justified, or a proposition with a probability of 0 is justified. Both horns of the dilemma are taken by DW to be absurd, leading them to conclude that *D* is not structural. The proof is understood by DW as showing that any formal solution to the paradox fails. I argue that DW's view that the proof is a reductio of *D* being structural is incorrect. I do this by providing reasons in favor of the claim that the source of the dilemma is the assumption that justification is aggregative (where a property *P* is aggregative iff, if $P\varphi$ and $P\psi$, then $\langle P\varphi$ Ù $P\psi\rangle$), rather than the assumption that *D* is structural.

undermine (i.e., Equal Likelihood), the door remains open to solving the lottery paradox with an account of justification formulated in purely probabilistic notions.

4. Non-Probabilistic Accounts of Why Lottery Propositions Lack Justification

In this section I will critically examine some non-probabilistic views of No-Justification. By non-probabilistic views I mean those views that explain the lack of justification of lottery propositions in terms not ultimately reducible to logical and probabilistic notions. I will discuss Igor Douven's pragmatic account, Jonathan J. Ichikawa's knowledge-first account, Christopher Kelp's virtue epistemology account, and Martin Smith's normic support account.

It is also important to note from the outset that it is not the explicit purpose of Douven, Ichikawa, and Kelp to explain why we do not know lottery propositions in terms of why we are not justified in believing them. It is interesting, however, to examine whether such accounts of No-Justification allow us to explain why we do not know lottery propositions.

1. Douven's Pragmatic Account

In "The lottery paradox and the pragmatics of belief" (2012), Igor Douven argues that lottery propositions fail to be rationally credible or epistemically justified because of Gricean considerations about what they *implicate* - in Grice's sense of *suggest, indicate, imply*.[38] He first provides a Gricean account of why lottery propositions cannot be warrantedly asserted, and then

[38] See Grice (1989: 42).

draws an analogy with rational belief. I will briefly reconstruct the above steps, and then present and evaluate his view on why lottery propositions are not epistemically justified.

Suppose you buy a lottery ticket, *t1*, and that Tim, a serious-looking man approaches you and in a neutral and seemingly sincere way says “I heard that you bought ticket *t1*. I have to tell you that *t1* won’t win.” He is seriously asserting it, so you ask yourself why Tim is telling you this. Douven considers different candidate explanations: i) Tim wants to convey that he possesses inside information about the outcome of the lottery that implies that *t1* won’t win, or ii) he wants to convey that he has inside information implying that *t1* is at least more likely to lose than what public information suggests, or iii) he wants to remind you of the low chances of *t1* winning, or maybe iv) he thinks that you may not realize how low your chances are and wants to direct your attention to it. According to Douven, Tim could have easily expressed himself more clearly if he wished to convey the ideas in (ii), (iii) or (iv), saying something like ‘your ticket is even less likely to win than you may think’ or ‘your ticket almost certainly won’t win’ - but he didn’t. Additionally, suppose that in following the Gricean maxims of Quality and of Manner, Tim is being cooperative -i.e., he tries to make a true, relevant contribution to the conversation- and in trying to be perspicuous he makes an effort to avoid ambiguities (Grice, 1989:27). This suggests that if he wished to convey the ideas in (ii), (iii) or (iv) and could have easily done it by expressing himself more clearly, he would have done so. So, it seems that at least one plausible explanation of why Tim asserted that *t1* won’t win is that he intended to convey that he has inside information implying that *t1* won’t win – that is, explanation (i). [39] With his assertion, Tim is raising suspicion,

[39] Douven also considers the objection that to convey the idea in explanation (i), Tim could have also easily expressed himself more clearly, saying ‘Your ticket certainly won’t win’ or ‘Your ticket is a sure loser’, among others. He dismisses it as an objection, appealing to the *weak* nature of the assertion’s implicature (360). If it were clearly the best explanation for Tim’s utterance, the implicature would be strong. For ways in which the implicature can vary in strength, see Douven (2012:354-355).

suggesting, or implying that he has such information. He is conversationally implicating that he has such information.

Admittedly, Tim’s utterance does not prompt you to fully believe that he has inside information. Douven holds that the relevant implicature is *weak*, in the sense that Tim’s assertion would make you believe only to *some* degree that he has inside information. Maybe (i) is not clearly the best explanation for Tim’s assertion, but there’s also no clearly better explanation than (i). However, since Tim is asserting a lottery proposition, and lottery propositions are about tickets in a fair lottery whose drawing is yet to occur, he does not in fact have inside information on the outcome of the lottery. In asserting that *t1* won’t win, Tim is *misleading* you into suspecting that he has inside information. But following Grice, a proposition P cannot be warrantedly asserted if, assuming that the speaker is being cooperative, asserting P would implicate a falsehood or would in some other way mislead the hearer (even if P is true). Thus, <*t1* won’t win> cannot be warrantedly asserted.

Douven argues that the above Gricean considerations should play a role in determining what propositions we are epistemically justified in believing, and not only in determining what can be warrantedly asserted. To understand their role, it is useful to think of adopting beliefs as storing them in our minds, and of storing and retrieving beliefs in our memory as a way of communicating between our different selves at different points in time. For example, you adopt the belief in *Q*:<Peggy’s car is plain blue> on the basis of having seen it, and you store it in your memory box. Even though *Q* entails *Q**:<Peggy’s car is blueish>, *Q** pragmatically implicates that her car is to some extent not blue, that is, it implicates not-*Q*.[40] Since you are trying to be cooperative and to avoid ambiguities when communicating with your future self who could

[40] It is not clear that Q entails Q*, since being blueish might entail being close to but not quite blue. For the sake of the argument, I grant Douven (2010:38) that Q entails Q*.

retrieve your beliefs, you would not store a belief that clearly suggests a falsehood (even more so, a falsehood that contradicts your evidence of having seen that Peggy's car is plain blue). Douven codifies this idea in a pragmatic principle dubbed 'Epistemic Hygienics', according to which it is not rational to adopt beliefs whose adoption risks misleading your future self. [41] If your future self happens to forget that you saw that the car is plain blue, in adopting a belief in *Q** you would risk misleading her into thinking something false (and contrary to your evidence). So, it is not rational to adopt a belief in *Q**.[42] Thus, just as a proposition *p* cannot be warrantedly asserted if asserting *p* would be in some way misleading, adopting a belief in *p* is not rational if adopting it would risk misleading our future selves.

Adopting the belief in the lottery proposition P1:<*t1* won't win> would be similar to adopting the belief in *Q**. Suppose you committed to memory the belief in P1. For Douven, this is much like writing a note to oneself without knowing in advance if one's future self will read it. Keeping this in mind, suppose further that at some point in the future when you have forgotten all about the relevant lottery, you come across your mental note P1, and you ask yourself why you have stored this particular note in your memory box. Assuming your past selves to be cooperative and trying to avoid ambiguities, you notice that at least one prima facie plausible explanation of why you stored P1 is that you had inside information implying that *t1* won't win, and that there does not seem to be a clearly better explanation than this one. This would be enough to suspect, even if only weakly, that you had inside information at the time you stored your belief. Thus, in committing P1 to memory you risked misleading your future self into falsely thinking that you had inside information. Since adopting a belief in P is irrational when it risks misleading our future

[41] For a defense of Epistemic Hygienics, see Douven (2010).
[42] This, of course, presupposes that our memory sometimes fails us and that your future self could forget that you saw that Peggy's car is plain blue. Douven welcomes this assumption (2012: 364-365).

selves into falsehoods, adopting your belief in P1 was irrational. Your belief in P1 is unhygienic, and therefore not epistemically justified.

1. Evaluation of Douven's Pragmatic Account

An initial worry with Douven's view of why we are not justified in believing lottery propositions is that there is a conceptual gap between his *pragmatic* principle against holding unhygienic beliefs (i.e., beliefs suggesting falsehoods, or contradicting the available evidence), on the one hand, and the *epistemic* justificatory status of such beliefs, on the other. However, the alleged gap is not as hard to fill. Let us understand the property of the belief that P of being epistemically justified as the property of the belief that P respecting the epistemic norms regarding doxastic states. Douven, then, can argue that lottery propositions fail to respect some important epistemic norms because of being unhygienic. For example, Douven might appeal to the norm that one should try to maximize the ratio of true beliefs to false beliefs, and argue that unhygienic beliefs produce the contrary effect; or he might appeal to the norm that one should not have beliefs contradicting one's evidence, and argue that unhygienic beliefs mislead us into accepting information contradicting the available evidence. Douven's explanation, however, is far from complete, since even though unhygienic beliefs have the potential to mislead us to a certain extent, this is different from showing that they indeed tend to mislead us to a significant extent, enough to disqualifying them epistemically. Moreover, whether the misleadingness of a belief epistemically disqualifies it is expected to vary with the specific context in which the belief is deployed, since the same occurs with any other pragmatic phenomena. However, Douven's view is prima facie workable as a way of supporting the claim that pragmatic factors might affect the justificatory status of a belief.

A more pressing concern with Douven's view is that it is difficult to make exact sense of it, because the most direct reading of it is conceptually problematic. Firstly, Douven states that a belief is not individuated coarsely grained, merely by its truth conditions, but finely grained, partly individuated by the information pragmatically conveyed by the assertion of the sentence semantically expressing its content (2012:364). Let us call this reading of Douven's view 'Douven's Encroachment of Content' or 'DEC' for short. DEC entails that if 'S_1' and 'S_2' are a pair of logically equivalent sentences commonly used to convey non-equivalent (or even incompatible) pieces of information, the belief reported by 'I believe that S_1' could be justified while the belief reported by 'I believe that S_2' is unjustified. For example, according to Douven, even though 'he was thirsty, but the milk was spoiled' is logically equivalent to 'the milk was spoiled, but he was thirsty', the beliefs reported with each sentence might vary in justificatory status since the former sentence conveys the piece of information that the person did not drink the milk while the latter sentence conveys the piece of information that he drank the milk, and one cannot be justified in believing both pieces of information. Given the widely accepted claim that a belief is a propositional attitude, i.e., a relation between an agent and a *proposition* (that is, a non-linguistic, mind-independent entity, different from the sentences expressing it, the assertions made by their use, and the mental representations and states subjects have towards it), DEC amounts to saying that the content of the belief *that S* is not merely the proposition semantically codified by 'S', but the proposition conveyed in the context in which 'S' is being employed.[43] Let us illustrate DEC with the next table. Let *P* be the proposition semantically codified by 'The milk

[43] A similar reading takes sentences to be the objects of beliefs, which is supported by Douven's own way of speaking. The resulting view seems to have the same problems as DEC, presented below, since the problematic information suggested by a belief would seem to be entailed by it, and it changes the belief under epistemic assessment with an irrelevant belief pragmatically expressed by the same sentence. Moreover, the view that the objects of beliefs are sentences faces clear problems when considering the trivial fact that the same belief can be reported with different sentences of the same or different languages.

was spoiled, but he was thirsty', *Q* be the proposition semantically codified by 'He was thirsty, but the milk was spoiled', and *R* be the proposition semantically codified by 'He drank the milk':

Sentence	Semantic Content	Proposition conveyed	Belief's content
'The milk was spoiled, but he was thirsty'	*P*&*Q*	*R*	*P*&*Q*&*R*
'He was thirsty, but the milk was spoiled'	*Q*&*P*	¬*R*	*Q*&*P*&¬*R*

Notice, however, that under DEC, holding a given belief not only *suggests* a false proposition, or a proposition contrary to the available evidence, but *entails* it. This is contrary to both Grice's conversational view on implicature and Douven's own explanation of why unhygienic beliefs are epistemically faulty, since the relation between the target belief and the problematic proposition being suggested by it is weaker than entailment, according to both. Moreover, DEC risks changing the belief that is the target of epistemic assessment to a different belief, pragmatically expressed by the same sentence, which is irrelevant to the topic under discussion. For example, if we are assessing a closure principle of justification with respect to the rule of commutation of the conjunction (If A is justified in believing P&Q, then A is justified in believing Q&P), DEC seems to entail the bizarre (and completely irrelevant) prescription that we should consider which propositions are conveyed by the sentences semantically codifying each conjunctive propositions in a given context to assess whether the principle holds in such context.

There is a similar worry to consider when applied to the case of lottery propositions. Suppose we ask if John is justified in believing that his lottery ticket will lose based solely on his evidence that the odds that it wins a fair lottery are extremely low. DEC recommends considering the epistemic status of the proposition that John's lottery ticket is going to lose *and John (or some possible interlocutor, perhaps) has inside information that the lottery ticket is going to lose*, which is completely irrelevant to the epistemic status of the former belief.

We can extract from the previous considerations that DEC is presumably not Douven's real position, and that the citation supporting it is rather an unfortunate formulation if it. In my opinion, the best way of understanding Douven's view is as a view on how beliefs are stored in one's memory. We can view memory as a box containing items in different compartments, where believing is one such compartment (along with doubting, desiring, and other mental attitudes). Thus, the belief in a proposition P would amount to storing a syntactic representation of P in the belief compartment. The items in each compartment are not propositions, since they are abstract entities that cannot be physically stored, but rather syntactic representations of beliefs. Under this view, how we store information is very relevant to how we process it, since if we use a misleading syntactic representation of a proposition, our memory might end up retrieving it in a context in which an unintended proposition ends up being expressed by it. This picture illustrates how the sentences used to express our beliefs would be constitutive of them (i.e., part of what individuates them) without being their content, and why it is relevant to consider the information suggested by them to epistemically assess them. Since even if a piece of information is not part of a belief's content, retrieving the belief from memory in a given communication context might produce beliefs in problematic propositions, in virtue of the syntactic item constitutive of the belief. However, the picture is rather cartoonish, and in order to consider it seriously, this understanding of Douven's account should be fully fleshed out, explaining what the syntactic representations stored in the brain are, how they convey wrong information in a context, what it is for a proposition to be "intended", and so on.

Douven's general idea of why being hygienic is significant for epistemic justification has merit. However, I believe that it ultimately fails in the case of lottery propositions, for two main reasons. The first is that Douven's explanation only works in conversational contexts, in which the intentions of the interlocutor and the evidence available to them is opaque to us. In such context, it makes sense to make abductive inferences regarding what she might have meant by her assertion that my lottery ticket is going to lose. But consider the situation from my interlocutor's perspective, who is convinced based on purely probabilistic evidence that my ticket is going to lose. Which false information is he risking convincing himself of in the future if he keeps the belief that my ticket is going to lose? Certainly, not the information that he has inside information regarding my lottery ticket being a loser, or anything similar, since no such information might ever play a role in the cognitive processes making him conclude, and maintain a belief of, the conclusion that my lottery ticket is going to lose. Would Douven be willing to accept that my interlocutor is indeed justified that my lottery ticket is a loser? I do not think so, since his explanation attempts to account for *all* cases of lottery propositions (at least, those involving normal human agents). The second reason is briefly considered and quickly dismissed by Douven himself (2012:364). Douven considers the objection that his theory entails that an agent with perfect memory is justified in believing lottery propositions (and, in general, unhygienic propositions), since she would not risk misleading herself into believing something false based on a faulty retrieving of her stored beliefs. Douven's response to the objection is simply to bite the bullet, claiming that his theory is only intended to account for the epistemic justificatory status of the beliefs of normal human agents. However, the arguments in section II of this paper supporting the claim that beliefs in lottery propositions are unjustified, are equally applicable to the case of agents with a perfect memory. Moreover, it is clear that such agents would not be in a position to know lottery propositions, regardless of their perfect memory, and it is reasonable to think that whatever correctly explains why they are not in such position to know is the same which correctly explains why we are not in such position to know, so Douven's refusal to extend his explanation to such ideal agents indicates that his proposal is not the ultimate explanation of why lottery propositions cannot be known.

2. Ichikawa's K-First Account

In 'Justification is Potential Knowledge'(2014), Jonathan Jenkins Ichikawa provides a knowledge-first account of doxastic justification that implies that lottery beliefs are not justified. His account aims to explain justified belief in terms of knowledge without identifying both:[44]

> **JPK** S's belief is justified iff there is a possible duplicate S′ of S with respect to a relevant set R of intrinsic states such that S′'s corresponding belief is knowledge. (189)

The notion of intrinsic state is considered to be an intuitive one, so Ichikawa treats it as primitive. However, the set R of relevant intrinsic states is restricted in various ways. In the first place, Ichikawa takes JPK to be 'mentalist' in the sense of Conee and Feldman (2001:2), i.e., it entails the claim that justification supervenes strongly on mental states. Because of this, R is taken to be a set of intrinsic mental states, excluding factual mental states (e.g., seeing, remembering, knowing), and mental states with contents individuated in an externalist way. The motivation for this restriction is to preserve intuitions about the notion of a justified belief, such as that there are Gettier-like cases of justified true beliefs that fall short of knowledge, and that the corresponding beliefs of my duplicates in all relevant intrinsic respects are as justified as mine, even if my

[44] Ichikawa's motivation for preserving a conceptual distinction between justification and knowledge is to agree with some of our epistemic intuitions, such as the intuitions that there are justified false beliefs and that there is a distinction between knowledge and justified true belief, exemplified by Gettier cases (2014: 188-9).

duplicates are in a twin-earth-like or brain-in-a-vat-like scenario. Moreover, R is not taken to be the total set of one's intrinsic mental states, since Ichikawa wants to make room for the possibility of *false* justified beliefs about a given intrinsic mental state (if I am justified by testimony in believing falsely that I was in pain while sedated during an operation, JPK entails that some duplicate of me in all relevant intrinsic states *knows* she was in the same type of pain, in which case we differ in one intrinsic state, i.e., she was in pain while I was not).

Applied to lottery beliefs, JPK has the consequence that, for every duplicate of an agent that has a statistical belief in a lottery proposition, she does not know the lottery proposition. According to Ichikawa, any intrinsic duplicate of an agent "will also […] be in the lottery situation, which seems to preclude knowledge". In other words, all the agent's duplicates will share the same intrinsic states as her, which include believing a lottery proposition based on mere statistical evidence. Since mere statistical evidence does not suffice for knowledge, the agent has no internal duplicates that know lottery propositions. Consequently, JPK yields that the agent's lottery belief is unjustified.

The above explanation gets things right extensionally, since (in accordance to the main intuition motivating this paper) we never know that a given lottery ticket is going to lose based on merely statistical evidence. More generally, JPK is attractive because it seems to preserve our epistemic intuitions about justification and knowledge as distinct concepts, as in Gettier-like cases. Nonetheless, JPK it cannot explain *that we do not know lottery propositions in virtue of not being justified in believing them*. For this explanatory purpose, as I argue next, JPK yields a mere restatement of the fact to be explained.

1. Evaluation of Ichikawa's K-First Account

Consider the central claim to Ichikawa's account that *every duplicate of an agent that has a statistically based belief in a lottery proposition does not know*. Let us now ask why every possible duplicate of an agent in such circumstances fails to know. If the response is merely that each of the corresponding beliefs of those duplicates are not justified, we have an uninformative response, since the idea that they are not justified is understood as explained by the idea that none of them are knowledge. This would make the considered explanation circular. If the response is that some state outside the relevant set R of intrinsic states is the culprit, then the reason why the agent does not know the lottery proposition is a non-justificatory one (since only the states in R are relevant to determining justification, according to JPK), so we are no longer explaining why the agent fails to know the lottery propositions (and statistical beliefs in general) in terms of her failing to be justified. Therefore, appealing to states outside R does not allow us to explain why the agent's belief fails to be knowledge by failing to be justified. Which other explanations can be given of the fact that every duplicate of an agent having a belief in a lottery proposition fails to know? Given Ichikawa's internalistic conception of justification (justification supervenes on a certain subset of mental intrinsic properties), it must be some of the intrinsic states shared by all and only the duplicates of the agent, but we are not told which one makes it impossible for any of them to know.

To summarize the problem, let J be the state of having only purely statistical evidence in favor of the belief in P. Plausibly, J is a relevant intrinsic state, since any possible duplicate of an agent with respect to J fails to know P. So, J belongs to R. According to Ichikawa's account, then, appealing to J is sufficient to account for why the agent fails to be justified in believing P, that is: since every duplicate of the agent with respect to J fails to know P, she is not in a position to know

P. However, the claim that any duplicate of an agent having J fails to know P is a restatement of the fact (which we wanted to explain in the first place) that beliefs in lottery propositions fail to be knowledge. Thus, Ichikawa's account does not provide an explanation of why beliefs in lottery propositions are not knowledge. Explaining *this* fact would call for a different, more informative, intrinsic state of the agent, but Ichikawa's account does not provide it, nor does it include a way to specify it. It gives the correct results only because it tacitly relies on the fact that statistical beliefs are unknown, which is the fact that we are trying to explain in the first place. Of course, the account is compatible with the existence of an informative intrinsic state different from J that accounts for why the agent does not know P, but then such intrinsic state is the one we should try to identify to get the desired explanation. Ichikawa's account is not doing the work.

3. Kelp's Virtue Epistemology Account

In "Lotteries and Justification" (2017), Christopher Kelp explains No-Justification with an account of justified believability (of propositions) built upon a virtue-epistemology (VE) and a knowledge-first framework. Within the VE framework, Kelp identifies justified belief with competent belief, where competently believing that *p* partly results from an exercise of a cognitive ability under a given set of conditions of shape and situation.[45] For example, we may form competent beliefs about our mental states by exercising introspection when we are in good shape by being awake, sober, and not distracted. Or we may form competent perceptual beliefs by exercising our faculty

[45] I am here narrowing the framework of VE to the standardly known category of virtue *reliabilism*, under which Kelp's virtue epistemology falls. Virtue reliabilists (e.g. Greco and Sosa) tend to identify intellectual virtues with cognitive faculties like perception, introspection, and memory, which are normally taken to reliably produce true beliefs. In contrast, virtue *responsibilists* (e.g. Code, Montmarquet, and Zagzebski) tend to identify intellectual virtues with character traits like intellectual openness, conscientiousness, and intellectual courage. On the division between virtue reliabilism and responsibilism, see Axtell (1997), Battaly (2015), and Fleisher (2017).

of vision when in good shape and situational conditions are met, e.g., there is normal lighting and other appropriate environmental circumstances. Yet, in contrast to VE classical accounts, competent belief is analyzed in terms of knowledge. Rather than featuring a cognitive ability whose exercise disposes us to form true beliefs (as abilities in VE classical accounts do), a competent belief features an ability whose exercise disposes us to form beliefs that qualify as knowledge.[46]

The abilities whose exercise can produce competent beliefs in inquiry are "abilities to know":

> One has an ability to know propositions in a range, *R*, relative to conditions *C* of shape and situation, just in case one has a grounded way of belief formation, *W*, such that using *W* in *C* disposes one to form beliefs that are about propositions within *R* and which qualify as knowledge. (2017a:9, 2017b:14)

According to this characterization, an ability to know ranges over a certain class of propositions. For example, we may have the ability to know a range of propositions R_D about deciduous trees but not a range of propositions R_E about evergreen trees. Such an ability amounts to having a way of forming beliefs that is *grounded*, that is, originated by human design or by natural evolution as a successful way of forming beliefs that qualify as knowledge. For example, visually perceiving

[46] See (Sosa, 2010) for more on conditions of shape and situation. Kelp's view takes our epistemic practice of inquiry as distinctively aiming at producing knowledge, rather than at producing mere true belief (2017a:9). Kelp considers the practice of inquiry about whether *p* (where *p* is any proposition subject to inquiry) to be a "simple goal-oriented practice" (SPG). SPGs are constituted by two types of particulars, moves and targets, and a designated relation between them. For example, the practice of archery ARC consists of shots taken from a distance (moves) aimed at discs (targets), where the designated relation is the hit relation. Similarly, the practice of inquiry involves ways of forming beliefs (moves), where the target of inquiry is any proposition *p*, if *p* is true (and $\neg p$, if *p* is false), and the designated relation is the knowledge relation between the belief and *p*. And, just as the practice of ARC is successful when it meets the goal of hitting the disk, the practice of inquiry is successful when it meets the goal of acquiring a belief that qualifies as knowledge. According to Kelp, traditionalist VEs can use the SPG framework, where the designated relation is a correspondence relation between the belief and *p* (when *p* is true), instead of the knowledge-relation (2016:14).

something under normal conditions of shape and situation is an ability that originated in our natural history as a successful way of producing beliefs that qualify as knowledge, which can range at least over propositions about medium-sized objects in our immediate surroundings.

Under Kelp's view, forming a *competent* belief in *p* amounts to i) exercising the ability to know *p* and ii) *p* falling within the range of target propositions relative to such ability. Without (ii), exercising an ability to know *p* is not enough to produce a competent belief, since an ability cannot provide knowledge of a set of propositions not belonging to its range.[47]

Kelp seems to understand a proposition *p* being justifiably believable as the idea that one is in a position to justifiably believe *p*, and since he identifies justified belief with competent belief, this notion amounts to being in a position to form a competent belief in *p*. In other words,

> Justified-Believability. A proposition *p* is justifiably believable for one if, and only if, one is in a position to believe *p* via an exercise of an ability to know propositions within range *R* and relative to *C* such that $p \in R$.(2017a:10)[48]

Along with the characterization of an ability to know, Justified-Believability allows Kelp to explain why lottery propositions are not justifiably believable for us. Consider a typical agent, John, who believes the proposition that his lottery ticket is not the winner —let us call it '*p*'— produced by his ability of inferring a proposition from purely statistical evidence showing that the

[47] I am restricting this presentation to competent (i.e., justified) belief, which Kelp distinguishes from apt belief (i.e., beliefs that are knowledge). Under his view, one may competently believe *p* without aptly believing it when conditions of situation are not met. For example, one might be in Fake-Barn County looking at the only real barn. In this scenario, the subject may competently form the belief that there is a barn in front of her, given that she exercises her ability to know barns relative to situational conditions SI_B in which what looks like a barn is in fact a real barn, and that the target proposition that she believes is within the range of her ability. However, she is not actually in SI_E (2016:15-6).

[48] Kelp identifies justifiably believing *p* with forming a competent belief in *p*. Also, Kelp seems to take as equivalent "*p* being justifiably believable" for someone and "being in a position to justifiably believe *p*". If this is correct, being in a position to justifiably believe *p* is equivalent to being in a position to form a competent belief in *p*. Since forming a competent belief in *p* consists in exercising an ability to know such that *p* falls within its range, being in a position to form a competent belief in *p* consists in being in a position to exercise said ability such that *p* falls within its range.

chance of that proposition being true is extremely high but lower than 1 (e.g., the chances of the ticket winning is one in a million, so *p*'s chances of being true are very high but not one). Let us call this ability 'PROB'. Kelp's explanation of why John's belief is not justified amounts to a story of the following type. For a belief to be justified is for it to be a competent belief, i.e., being produced by an ability-to-know propositions of a certain range. As previously mentioned, an ability counts as an ability-to-know propositions in range *R* and in circumstances *C* if and only if such ability was originated as a successful way of producing knowledgeable beliefs. PROB, however, was not originated (either by design or by natural evolution) as a successful way of producing knowledgeable beliefs in lottery propositions. PROB, of course, constitutes a disposition to produce true beliefs (it would produce a false belief only in an insignificant number or cases), but that does not make it an ability *to know*, just as shooting an arrow without any technique towards a very large target is the result of mere chance rather than of a genuine ability, even if the arrows hit the target most of the time.

1. Evaluation of Kelp's Virtue Epistemology Account

A full evaluation of this proposal requires determining if Kelp's account of justification delivers all and only correct results, and this is beyond the scope of this paper. In favor of the proposal, we can say that it clearly delivers the right results for the case of beliefs in lottery propositions, since all of them are formed by PROB (or similar ways of acquiring beliefs in the usual circumstances in which the only available evidence is statistical). This is confirmed by the intuition, characterized at the beginning of the paper, that in typical cases we are not in a position to know lottery propositions.

Nevertheless, Kelp's proposal is unsatisfactory in the present context in which we are trying to explain the intuition that we do not know lottery propositions by means of an account of why beliefs in them are unjustified. In this context, answering that those beliefs are not knowledge because they fail to be justified, and that they fail to be justified because they are not produced by an ability-to-know (i.e., originated as an apt way of successfully gaining knowledge) does not provide the explanation we want. The answer does not seem more informative than claiming that we do not know lottery propositions because we do not have an ability to know them. Moreover, it is perfectly reasonable to expect that questions like "Why is perception-in-normal-circumstances a way of successfully gaining knowledge but not PROB?" have explanatory answers, just as it is reasonable to expect that it can be explained why some abilities are successful at producing certain goals unlike other similar abilities. However, Kelp's account seems unfit to provide such an answer. This is even clearer if we accept that there are cases of knowledge gained by non-deductive inferences, since then we can ask an even more pressing question: Why does purely probabilistic reasoning not yield inductive knowledge, even if the strength of such reasoning is as high as we would expect of any other non-deductive reasoning that does yield knowledge (even after restricting our attention to cases in which all the premises are known)? Let us contrast, for example, the following two cases (Harman, 1968:166):

Lottery Testimony

1. The newspaper announced that the lottery winning ticket number is 002323, which is not my ticket number,

Therefore:

2. My lottery ticket is a loser.

Lottery Probability

1. There is a 99.999 percent chance that my lottery ticket does not win,

Therefore:

2. My lottery ticket is a loser.

Lottery Testimony is non-deductive, since it is compatible with the newspaper misprinting the ticket number or with making a mistake while reading the newspaper, and even if the likelihood of that unlikely scenario were higher than .0001, we would consider Lottery Testimony a case of knowledge under normal circumstances. On the other hand, we would not consider Lottery Probability a case of knowledge under normal circumstances. Saying that the first form of reasoning was originated as a way of producing knowledge while the latter is not leaves it a complete mystery why one is apt for such a goal while the other is not.

One might think that the present objection against Kelp's proposal tacitly assumes that knowledge has a reductive analysis, but this is incorrect unless it is accepted that only reductive accounts of knowledge can provide explanatory answers to reasonable questions about particular cases. Although it is unreasonable to expect that a non-reductive account of knowledge provides an explanatory answer to why every putative case is, or fails to be, knowledge, it is unclear why the case of beliefs in lottery propositions should lack explanatory answers. If it should, it would be rather surprising, and Kelp's account does not provide even an inkling of why that would be the case.

4. Smith's Normic Support Account

Martin Smith (2016, 2010) argues that lottery propositions are not justified by introducing a condition for justification that is not reducible to relations of probability or likelihood, and arguing that lottery propositions do not meet that condition. Smith's condition is motivated by the intuitions that justified beliefs are candidates for knowledge, and that the evidence supporting a belief that is a candidate for knowledge is so good that it requires providing some explanation when the belief ends up being false. For example, suppose you leave the living room and come back a minute later. Before re-entering the room, you hold the justified belief that the furniture is still in place, having seen it a minute ago and having no reason to think it has been removed. To your surprise, you come back to an empty room. In this case, your belief being false warrants an explanation – maybe you've been subjected to a highly-elaborated prank, or maybe you are hallucinating. Being pranked or hallucinating are not normal circumstances, and the error in your belief seems independent of your intellectual responsibilities, e.g., it seems to be due to environmental conditions (as in the first case) or cognitive malfunctions (as in the second case). In any case, an explanation that accounts for why your belief is false in this case is needed, since your belief was a very good candidate for knowledge, given your evidence.

To accommodate the above intuitions and their exemplifications across numerous candidates for knowledge, the condition on justification proposed by Smith focuses on a kind of support between propositions and their evidence, which he dubs 'normic support':

> *Normic Support*) A body of evidence E normically supports a proposition P just in case the circumstance in which E is true and P is false requires more explanation than the circumstance in which E is true and P is true. (2016:40)[49]

'Normic' alludes to a comparison between how normal the circumstance is in which both the target proposition (P) and the relevant evidence (E) are true, and how normal the circumstance is in which the target proposition is false (¬P) while the same evidence remains true. Intuitively, a circumstance in which the furniture is still in the living room after having seen it a minute ago is

[49] Another characterization that Smith provides in terms of normalcy of worlds is as follows: A body of evidence E normically supports a proposition P just in case the most normal worlds in which E is true and P is false are less normal than the most normal worlds in which E is true and P is true. (2010:16-7).

more normal than a circumstance in which the furniture is not in the room a minute after having seen it. In this respect, *Normic Support* assumes a link between how normal a circumstance in which P is true is (given that E is true), and how much of an explanation would be required for a circumstance in which P is false (while E remained to be true). In Smith's words, "normal conditions require less explanation than abnormal conditions do" (2016:39). Think of the circumstance in which the furniture is no longer in the living room despite you having seen it there a minute ago. Intuitively, a circumstance c_1 in which the room floor collapsed in a matter of seconds due to moist and termite damage is more normal and requires less explanation than a circumstance c_2 in which an evil demon emptied out the room in a matter of seconds (after all, how did he manage to do that? did it involve magic?). Nonetheless, c_1 is still less normal and requires more explanation than a circumstance in which the furniture is still in the room (since the latter does not require any explanation). Thus, rather than intending to capture statistical frequency -as in uses of 'normal' like 'Rain is normal this time of year'-, normalcy seems to be characterized along the following lines:

> *Normalcy*) E&P is more normal than E&¬P if and only if E&¬P requires more explanation than E&P.[50]

Bearing in mind Smith's characterization of normic support, his proposed conception of justification is as follows:

> *Normic Justification*) For any proposition P supported by a body of evidence E, P is justified for S only if E normically supports P

[50] For more examples where normalcy and frequency come apart, see Smith (2016:39-40).

Normic Justification renders lottery propositions unjustified. First, let us remember that the relevant understanding of normalcy is not intended to capture statistical frequency. If it were, it would be more normal (i.e., more frequent) to lose the lottery than to win it. However, a lottery proposition can be true more frequently than not despite not being normically supported by the evidence. This is the case of any lottery proposition. A circumstance in which the proposition p_1:<ticket #1 is a losing ticket> is false and the probabilistic evidence *e* supporting p_1 is true, is equally normal as the circumstance in which p_1 is false and *e* is true. If ticket #1 loses, there's no need to explain why it lost; and if it wins, although this fact may surprise the ticket holder, there is also no demand for an appropriate explanation of why it won. In other words, the outcome in which a given lottery ticket loses is explanatorily on a par with the outcome in which it wins, so no proposition expressing either outcome is normically supported by mere statistical evidence. Therefore, no belief in a lottery proposition is justified on the basis of statistical evidence alone.[51]

1. Evaluation of Smith's Normic Support Account

Smith's support of his view that lottery propositions are unjustified in the envisaged cases relies on his claim that lottery propositions are not normically supported by purely probabilistic

[51] A similar view is proposed by Dana Nelkin (2000), according to which being justified in believing *p* requires that, on reflection, we are able to rationally postulate some causal or explanatory connection between our belief in *p* and the fact that makes *p* true. For example, I'm justified in believing that I will drink coffee in a moment. The reason for my belief is that I intend to do so. Under Nelkin's view, such intention both explains why I have this belief and the fact that I will drink coffee shortly. But in the lottery case there is no causal or explanatory connection between my belief that my ticket will lose and the fact that it will (when it loses). Our evidence is not a common cause nor a common explanation of our lottery belief and the fact that our ticket loses, nor is there any other causal or explanatory connection that we can postulate under reflection. According to Nelkin, an indication that there is no such connection is that if we were to find out that our belief is false (that is, that we won the lottery), we would not need to reject or revisit the evidence for our belief. This indication seems to echo Smith's claim that it would not be abnormal if our lottery belief were false, in the sense that there would be no need to seek an explanation for why it is false – for example, the explanation that something is amiss with the evidence for our lottery belief. If it turns out that I do not drink coffee shortly, this requires an explanation. Maybe I intended to drink it but was interrupted by an important call before drinking it. Then however good my evidence was, it did not guarantee that I drank the coffee.

evidence. However, Smith's motivation for this claim is unclear. One line of reasoning that he could be taking to support it is the following:

1. For any *p* solely supported by probabilistic evidence *e*, the circumstance in which $e \& \neg p$ does not require more of an explanation than the circumstance in which $e \& p$.
2. Every lottery proposition p_x is solely supported by probabilistic evidence.

Therefore,

3. For any lottery proposition p_x and every probabilistic evidence e' such that p_x is solely supported by e', the circumstance in which $e' \& \neg p_x$ does not require more of an explanation than $e' \& p_x$.

Given *Normic Support*, it follows from (3) that lottery propositions are not normically supported. That Smith would accept the previous line of reasoning in favor of (3) is suggested by Smith's implicit commitment to (1) and (2), from which (3) logically follows.[52] However, as argued by Blome-Tillmann, there are counterexamples to (1):

> "Consider the bark beetle. Interestingly, the bark beetle isn't threatened by climate change, even though 98% of insects are. This is surprising and calls out for an explanation. Why isn't the bark beetle threatened by climate change, given that almost all other insects are?" (568)

The statistical evidence that 98% of insects are threatened by climate change supports the false proposition that the bark beetle is so threatened. Yet, the circumstance in which 98% of insects are threatened by climate change & the bark beetle is not *does require* more of an explanation than

[52] Michael Blome-Tillmann (2020) argues that Smith's views plausibly generalize into (1) (where (1) is equivalent to Blome-Tillmann's (G') principle (2020:568)).

the circumstance in which 98% of insects are threatened by climate change and the bark beetle is so threatened. Therefore, given *Normic Support*, the evidence that 98% of insects are threatened by climate change normically supports that the bark beetle is threatened by climate change. The bark beetle case is thus a case in which a proposition is normically supported by mere probabilistic evidence. In this way (1) is false, undermining the reasoning for (3), which motivated the claim that lottery propositions are not normically supported.

It is unclear that the above considerations should be a cause of concern for Smith, whose support for (3) might be merely the fact that what (3) describes is intuitively true. No additional motivation beyond such intuitive appeal seems to be required. Moreover, Blome-Tillmann's own counterexample relies on the similar intuition that the proposition that bark beetles are not affected by climate change requires more explanation given the statistical evidence that 98% of insects, and no additional reason supporting it seems to be required. If this is the case, Smith can rely directly on his intuitions about whether a given circumstance requires more explanation than another circumstance.

5. Concluding Remarks

We have examined two types of view to account for why we do not know lottery propositions in terms of why we are not justified in believing them: probabilistic and non-probabilistic views. Although probabilistic views seemed problematic, I have argued that developing a view of this sort remains a theoretical possibility. We have also examined some salient non-probabilistic explanations. As we have seen, all but one of them seemed to face important problems making them unsuitable to play the explanatory role sought throughout this paper. The position that seems to best fulfill this role is Smith's view. It remains to be determined whether his view relies on the general claim that no proposition supported by mere probabilistic evidence is normically supported, which is subject to counterexamples, or whether it relies on our intuition that lottery propositions lack normic support.

References

Blome-Tillmann, M. (2020). Statistical evidence, normalcy, and the gatecrasher paradox. *Mind*, *129*(514), 563–578.

Douven, Igor (2002). A new solution to the paradoxes of rational acceptability. British Journal for the Philosophy of Science 53 (3):391-410.

(2010). The pragmatics of belief. Journal of Pragmatics, 42(1), 35–47.

(2012). The lottery paradox and the pragmatics of belief. *Dialectica*, *66*(3), 351–373.

Feldman, Richard & Conee, Earl (2001). Internalism Defended. *American Philosophical Quarterly* 38 (1):1 - 18.

Foley, R. (1987). *The theory of epistemic rationality*. Cambridge, MA: Harvard University Press.

Gilbert Harman. (1968) Knowledge, Inference, and Explanation. *American Philosophical Quarterly*, *5*, 164-173.

Ichikawa, J. J. (2014). Justification is potential knowledge. *Canadian Journal of Philosophy*, *44*(2), 184–206.

Kelp, C. (2014). Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51, 227–232.

(2016). Justified Belief: Knowledge First-Style. *Philosophy and Phenomenological Research*, *93*(1), 79–100.

(2017a). Lotteries and justification. *Synthese*, 194(4), 1233-1244.

(2017b). Knowledge first virtue epistemology. *Knowledge first: Approaches in epistemology and mind*, 223-45.

(2019). How to be a reliabilist. *Philosophy and Phenomenological Research*, 98(2), 346-374.

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

McGlynn, Aidan (2012). Justification as 'Would-Be' Knowledge. Episteme 9 (4):361-376.

Moser, P. K., & Tlumak, J. (1985). Two paradoxes of rational acceptance. *Erkenntnis*, *23*(2), 127–141.

Nelkin, D. K. (2000). The lottery paradox, knowledge, and rationality. *The Philosophical Review*, *109*(3), 373–409.

Ochoa, A. (2021). The Aggregativity of Justification and De-Generalizing the Lottery Paradox. [Unpublished Manuscript]

Pollock, J. L. (1995). *Cognitive carpentry: A blueprint for how to build a person*. Mit Press.

Ryan, S. (1996). The epistemic virtues of consistency. *Synthese*, *109*(2), 121–141.

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

(2016). *Between Probability and Certainty: What Justifies Belief Oxford University Press, 2016, 213 pages*.

Sosa, E. (2010). How Competence Matters in Epistemology. *Philosophical Perspectives*, 24, 465-475

* Section 3's integration into the dissertation project

Section 2's conclusion that epistemic luck as unsafe true belief cannot explain why we fail to know lottery propositions, coupled with the fact that safety is widely agreed to be the best account of epistemic luck, motivated an examination of whether lottery beliefs are justified. But some motivation for this project was already present for theorists that wished to solve the lottery paradox without rejecting the *aggregativity* of justification (that is, that justification is a property that a conjunction has if all of its conjuncts have it). This is because, as we'll see in Section 4, to solve the lottery paradox one must either reject the aggregativity of justification or the premise that lottery propositions are justified. And, as stated by Douven and Williamson, many theorists believe that denying the aggregativity of justification is "too drastic" (2006:758), for aggregativity of justification is just a form of epistemic multi-premise closure and theorists tend to avoid denying epistemic closure.

Initially, probabilistic accounts of justification were proposed to explain how lottery propositions had a given justification-defeating structural property *D* preventing them from being justified. But Douven and Williamson's (2006) proof closed the door to all probabilistic accounts by arguing that if there is a defeating property, it *cannot* be structural. Naturally, non-probabilistic accounts arose to provide an explanation of why lottery propositions are unjustified by proposing non-structural defeating properties that lottery propositions could have. As seen in Section 3.4, however, all examined views except for Smith's (2010) are unsuitable to ultimately explain why we do not know lottery propositions. If our options of non-probabilistic views were limited before, now they are even more limited.

The previous result motivates Section 3.3.4, together with section 4 of this dissertation. They together reopen the door to probabilistic accounts of justification to explain why we fail to know lottery propositions. Section 3.3.4 shows how DW's argument against Ryan's (1996) and Douven's (2002) probabilistic accounts does not extend to very similar probabilistic accounts, while Section 4 restores the possibility of lottery propositions possessing a justification-defeating structural property.

4. The Aggregativity of Justification and De-Generalizing the Lottery Paradox

1. Introduction

Solving the lottery paradox about justification[53] can serve as a criterion to choose among alternative theories of justification, if not as a condition of adequacy for such theories. A popular way to deal with this paradox is to find a property shared by all competing lottery propositions (yet not by justified ordinary propositions) that precludes them from being justified, despite their high probability (given the evidence). Such property is then used to restrict the principle called 'Sufficiency', which operates in the paradox saying that if a proposition is highly likely given the evidence, then it is justified. Sufficiency is restricted by adding to its antecedent the condition that the relevant proposition does not possess the problematic property shared by lottery propositions. Igor Douven and Timothy Williamson ('DW' for short) present a proof (henceforth, 'the proof')[54] which they take to show that if such property exists, it cannot be a 'structural' property (Douven & Williamson, 2006), that is, a property fully expressible by means of logical and probabilistic notions. [55] Any such solution to the paradox is labeled 'formal' by DW, so they take the proof to

[53] This paradox is originally presented by Henry E. Kyburg (1961). It is presented as a paradox about the notion of "rationality" (what Igor Douven and Timothy Williamson call 'rational acceptability', 2006), which is equivalent to the notion of justification in the present context (however, see Sylvan (Sylvan, 2014) for arguments against the identification of rationality and justification).
[54] DW formulate both the proof and the lottery paradox in terms of the notion of rational acceptability rather than in terms of the notion of justification. This is unproblematic in the present context in which both notions are typically used interchangeably.
[55] More rigorously, a structural property is defined by DW as "any property P such that for any proposition φ and any automorphism of propositions *f*, φ has P iff *f* (φ) has P" (p.762), where the function *f* satifies the following conditions: for every x and every y, f(x & y) = f(x) & f(y), f(¬x) = ¬f(x) and Pr(x)=Pr(f(x)) (pp.761-762). More briefly, DW's automorphism is a type of indiscernibility with respect to probability and basic logical properties. Thus, their structural property is one that logically supervenes on probability and basic logical properties and is ultimately expressible by them.

show that any formal solution is doomed to fail. The proof constitutes a very important development, as many of the proposed solutions that take beliefs in lottery propositions as unjustified are formal.[56] It relies on the assumption called 'Aggregativity', which says that justification is an aggregative property (i.e., a property that a conjunction has if all its conjuncts have it). Then, it concludes by reductio that no structural property can be used to solve the lottery paradox by restricting Sufficiency. In this paper I give two independent arguments for thinking that Aggregativity is incorrect, so if my arguments are cogent, the proof is better understood as a reductio of Aggregativity than as a reductio of there being formal solutions to the lottery paradox.

The proof is presented with the lottery paradox about the notion of justification (rather than the notion of knowledge) as its background,[57] which arises from the following assumptions:

Sufficiency: If $\Pr(\varphi)|\mathrm{E} > t$, then φ is justified (for a given t such that $.5 \leq t < 1$),

Aggregativity: Justification is aggregative, where a property P is aggregative iff, if $P\varphi$ and $P\psi$, then $P\langle\varphi \wedge \psi\rangle$

Non-Zero Probability: If φ is justified, $\Pr(\varphi)|\mathrm{E} \neq 0$.[58]

Taken together, these assumptions lead to a paradox. Consider a lottery with one winning ticket and n tickets such that $1 > (n-1)/n > t$. Now, for each member x of the set L of propositions

[56] See, for example, Douven (2002), Pollock (1995), and Ryan (1996).

[57] It is a matter of contention whether and how the lottery paradox of justification and the lottery paradox of knowledge are related. Although both paradoxes are structurally identical, and this makes it plausible that they share a solution (Nelkin, 2000), it might be argued that they differ in solution (for example, Williamson (2000) and Pritchard (2007) treat beliefs in lottery propositions as falling short of knowledge for being unsafe, but they do not consider them unjustified), or that they differ in their status as genuine paradoxes (Bondy, 2013). This paper remains neutral with respect to how that contended issue is resolved.

[58] *Sufficiency*, *Aggregativity*, and *Non-Zero Probability* should be considered assumptions that concern, rather than absolute probabilities, the probability of propositions relative to the available evidence. In the example following the characterization of such principles, the available evidence consists in the lottery having only one winning ticket and n tickets such that $1 > (n-1)/n > t$.

of the form ⟨Ticket #*i* of lottery *l* is a loser⟩, Pr(*x*) = (*n* -1)/*n*. So, given Sufficiency, each member of *L* is justified, and given Aggregativity, the conjunction of all members of *L* is justified too. But the probability of such a conjunction is 0, so it is not justified given Non-Zero Probability. Contradiction!

Before continuing, some comments on how 'justification' is employed in this paper. Generally, the term designates the normative status of certain types of items given certain normative criteria, and in epistemology it is specifically used to designate the *epistemic* status of either propositions ("propositional justification") or of doxastic states ("doxastic justification"). My use of the term here is narrower for the following reasons. First, the relevant items in the present context are lottery propositions, which are justified inferentially. Secondly, the evidence on the basis of which lottery propositions are justified only consists of statistical information determining their probability, e.g., the number of lottery tickets, the number of winning tickets, the facts that the lottery is fair, and that its results have not been revealed. Thus, throughout this paper, I restrict the use of 'justification' to designate the status of being (epistemically) inferentially justified, rather than of being justified *simpliciter*. The reasons for this restriction will be made clearer further ahead.[59]

Let us now return to the paradox. On a superficial examination, one might suspect that Sufficiency is the culprit and should therefore be rejected to solve the paradox. However, according to DW, the next weakened version of Sufficiency remains plausible even if Sufficiency is rejected

[59] Inferential justification is a type of *doxastic justification* in which *the belief* in a proposition *p* is justified for an agent S *based on* evidence available to S from which it is epistemically admissible for S to infer *p*. This notion is closely tied to a corresponding notion of *propositional justification* in which a proposition *p* is justified for an agent S when it is epistemically admissible for S to infer *p* from the evidence available to S *even if S does not actually believe p based on such evidence*. Given this close relation between both notions, the following discussion applies equally to both.

(p.758). Where t is a given rational number such that $.5 \leq t < 1$, φ any given proposition, and D is a property that neutralizes (in some way) the justification of a proposition if possessed by it, [60] we have the next principle:

Restricted Sufficiency: If $Pr(\varphi)|E > t$ and φ does not have D, then φ is justified.

Aggregativity, Non-Zero Probability and Restricted Sufficiency are the main assumptions of the proof, which I here characterize in its most salient features. First, DW proved two important results:

i) If D is structural, then the property of being a φ such that $Pr(\varphi)|E > t$ and $\neg D\varphi$ (let us call this property 'undefeated warrant') is structural too.[61]
ii) If a property P is structural and sufficient for having a property Q, and Q is aggregative, then if $Pr(\varphi)|E < 1$ and φ has P, then $\bot$ (any proposition with a probability of 0) has Q (p.762).

According to (i), if D is structural, then undefeated warrant is also structural. According to Restricted Sufficiency, undefeated warrant is sufficient for justification. Since, according to Aggregativity, justification is aggregative, it follows by (ii) that if D is structural, then if φ has undefeated warrant and $Pr(\varphi)|E < 1$, then $\bot$ is justified. Given these results (DW claim) a dilemma emerges: if D is structural, either Restricted Sufficiency amounts to the trivial claim that having a probability of 1 is sufficient to be justified, or a proposition with a probability of 0 is justified.

[60] DW call *D* a defeater (p.758). I prefer to avoid such denomination to avoid confusion, since that term is oftentimes reserved, by defeasibility theories, for true propositions that are in a certain logical "defeating" relation to the justifying evidence for a belief that is otherwise a candidate for knowledge.
[61] Strictly speaking, they proved the proposition that "Any predicate defined purely in terms of structural predicates by means of the Boolean operators and quantification is structural" (p.764). Since having a probability above t is structural, (i) follows from such proposition.

Both horns of the dilemma are taken by DW to be absurd, leading them to conclude that *D* is not structural.

Following DW, I label ‘formal’ any solution to the lottery paradox that understands *D* as structural. The proof is understood by DW as showing that any formal solution to the paradox fails. In contrast, I here argue that DW’s view that the proof is a reductio of D being structural is incorrect. I do this by providing reasons in favor of the claim that the source of the dilemma presented by DW is the assumption that justification is aggregative (that is, Aggregativity) rather than the assumption that *D* is structural.

Here is how I proceed in the rest of the paper. In section 2, I motivate two important principles of justification, based on which I present an argument against Aggregativity. In section 3, I examine and reject some objections to such argument. In section 4, I present a form of the lottery paradox that is more basic than the meta-reasoning about justification discussed by DW. I argue that Aggregativity obstructs solving such basic form of the paradox, and that there is a plausible restriction of Aggregativity that allows a solution to it. In the light of the arguments presented in sections 2 and 4, I conclude that the proof is better understood as a reductio of Aggregativity.

2. The Argument Against the Aggregativity of Justification

When we think of justification, it is clear that inasmuch as it is the output of a function of the probability of a proposition given the relevant available evidence, such function is not primitive.

The justification of a proposition is *primarily* a function of the weighting of the evidence in favor and against it, according to epistemic norms. Thus, whatever the exact relation between probability and justification is, probability should be part of the weighting of evidence for or against a given proposition if it is to play a central role in our understanding of its justificatory status.

And it is indeed very plausible that probability plays a central role in understanding of the notion of justification, since it plays a central role in our understanding of the notion of a proposition being *supported* by the available evidence (arguably, the notion of support is *the* most central component of the relevant notion of justification). This intuition is captured as follows:

> *Simple Argument*
>
> If a given proposition is supported by some available evidence, it can be derived, either deductively or non-deductively, from such evidence. And the fact that a proposition can be non-deductively derived from some available evidence requires that its probability given such evidence is high enough (otherwise the derivation would not be a strong one, which is a necessary condition for it to support its conclusion). Thus, the probability of a proposition given the evidence should be high enough in order for it to be supported by such evidence. Since the justification of a proposition requires that such proposition is supported by the available evidence, it follows that a justified proposition should have a high enough probability given the available evidence.

I find this argument to be very intuitive. A closer look at it allows us to identify general principles concerning the notion of justification. Let E be the set of propositions consisting in all the pieces of available evidence, and of nothing else. The argument involves the following intuitive principles:

a. If j is justified, then E supports φ.

b. If E supports φ, then a derivation of φ from E is strong to a high enough degree (i.e., above certain threshold).

Let us remember that 'justification' is used throughout the paper to designate the notion of *inferential justification* rather than the notion of justification *simpliciter* (and similarly for related terminology). This is important because if principles Aggregativity and (b) are understood in terms of the notion of justification *simpliciter*, they are controversial (they amount to a sort of evidentialist account of justification) and might have plausible counterexamples (regarding "basic" propositions, whose justification is non-inferential). However, they seem to be obviously true when understood in terms of the notion of inferential justification. Given that lottery propositions are justified inferentially based on purely statistical evidence, it is safe to identify the relevant notion of justification with inferential justification.

Before presenting my argument against aggregativity, it is useful to make the following set of definitions and stipulations:

$Stip_1$. Pr(φ|E) = the degree to which the derivation of φ from E is strong.

$Stip_2$. Pr(φ) = Pr(φ|E).[62]

$Stip_3$. *t* is a degree such that φ is supported by E only if the degree of strength of a derivation of φ from E is above *t* (that is, Pr(φ|E) > *t*, according to $Stip_1$).

Def_1. φ is sufficiently probable = Pr(φ|E) > *t*.

[62] This stipulation should be read as an abbreviation of Pr(φ) in the context of discussing the probability of lottery propositions given the evidence, rather than as a general truth about the absolute probability of lottery propositions.

Def_2. φ is supported = φ is supported by E.

The rationale for $Stip_1$ and $Stip_2$ is to provide an epistemic interpretation of probability,[63] required if probability is to have a role in understanding the notion of justification. If, for example, probability were understood as a non-epistemic physical measure of some sort (e.g. a physical propensity or disposition) which is not reflected in the available evidence, it would be utterly mysterious why that measure is supposed to be conceptually connected with the notion of (epistemic!) justification.[64]

In contrast, $Stip_1$ and $Stip_2$ allow us to interpret probability as measuring the degree of support a proposition φ has given E, which corresponds to the degree of strength of a derivation of φ from E. $Stip_1$ identifies the degree of strength of the derivation of φ from a set of propositions Γ with the conditional probability of φ given Γ. This concurs with the intuitive idea that if the derivation of φ from Γ is deductively correct (that is, valid), then $Pr(\varphi|\Gamma) = 1$, and if it is non-deductively correct (e.g., if it is an inductively strong derivation), then $Pr(\varphi|\Gamma)$ is high but possibly below 1. This interpretation takes as basic the notion of conditional probability (in its epistemic sense) and defines with it the notion of absolute probability, as captured by $Stip_2$. As a last step, $Stip_3$ stipulates *t* to be a threshold for a derivation being strong enough for providing support. The

[63] This interpretation is based on the notion of probability Carnap used in his inductive logic (Carnap & Jeffrey, 1980), as the degree of confirmation by the evidence a proposition has.

[64] If Pr(φ) and Pr(φ|E) are priors, $Stip_2$ entails that E does not support φ, according to a Bayesian analysis of the notion of support or confirmation (that is, E supports φ iff $Pr(\varphi|E)_i > Pr(\varphi)_i$, where the subindex 'i' denotes that the probability in question is a prior (Talbott, 2016)). Since we are assuming that E supports φ, it follows that either we should understand Pr(φ|E) as a prior and Pr(φ) as a posterior in $Stip_2$, or that the argument proposed here is inconsistent with Bayesian accounts. This, however, is puzzling when φ is a proposition supported only on the basis of statistical evidence, because in that case E would only include information determining Pr(φ) (e.g. the number of tickets of the relevant lottery, the proposition that the lottery is fair, that there is only one winner, etc.), and there might be no probability assignment to φ prior to Pr(φ).

role of $Stip_3$ is to allow us to express the intuitive idea that the strength of a derivation is high enough in terms of a probabilistic value (or range of values).

Finally, Def_1 and Def_2 are mere abbreviations that simplify the presentation and discussion of the argument against the aggregativity of justification, which we can now formulate as follows:

I. If φ is justified, then φ is supported.

II. If j is supported, then φ is sufficiently probable.

III. If justification is aggregative, then support is aggregative.

IV. If support is aggregative, then sufficient probability is aggregative.

However,

V. Sufficient probability is not aggregative.

Therefore:

VI. Justification is not aggregative.

Let us consider the motivation for each premise. (I) and (II) are brief reformulations of the principles of justification presented at the beginning of this section, which I simply assume throughout this paper given their intuitive appeal. These two principles do not really function as premises of the main argument, but they are used for justifying the premises, as I show next.

(III) can be shown to be true by a reductio. Suppose that justification is aggregative but support is not. Suppose also that j and y are both supported, but φ∧ψ is not supported (which can happen if support is not aggregative, as we assumed). Suppose further that φ and ψ satisfy every necessary condition for justification in addition to being supported. Then, φ and ψ are both justified. And, given the aggregativity of justification, φ∧ψ is also justified. However, given (I),

$\varphi \wedge \psi$ is not justified, since it is not supported. Contradiction! Therefore, if justification is aggregative, so is support.

(IV) can be argued for in the same way as (III). Suppose that support is aggregative but sufficient probability is not. Suppose also that φ and ψ are both sufficiently probable, but $\varphi \wedge \psi$ is not sufficiently probable (which can happen if sufficient probability is not aggregative, as we assumed). Suppose further that φ and ψ satisfy every necessary condition for support in addition to being sufficiently probable. Then, φ and ψ are both supported, and given the aggregativity of support, $\varphi \wedge \psi$ is supported too. However, given (II), $\varphi \wedge \psi$ is not supported, since it is not sufficiently probable. Contradiction! Therefore, if support is aggregative, so is sufficient probability.

Finally, let us argue in favor of (V). First, provided that i) $0 < Pr(\varphi) < 1$ and ii) $Pr(\psi|\varphi) < 1$, it follows that $Pr(\varphi \wedge \psi) < Pr(\varphi)$.[65] Let Γ be a set of propositions such that any member of Γ satisfies (i) and any ordered pair of two members of Γ satisfy (ii). In that case, if $|\Gamma|$ is large enough, the probability of the conjunction of all the members of Γ might end up being equal to or below t even though each member of Γ has a probability above t. For example, consider the set $\{\varphi_1, \varphi_2\}$ such that $Pr(\varphi_1) = .7$, $Pr(\varphi_2) = .7$, and $Pr(\varphi_1|\varphi_2) = Pr(\varphi_1)$, and suppose that $t = .5$. Even though $Pr(\varphi_1) = Pr(\varphi_2) > t$, in this case $Pr(\varphi_1 \wedge \varphi_2) = .49 < t$.[66] Of course, there are countless models of

[65] The reason is simply that $Pr(\varphi \wedge \psi)$ is the multiplication of two fractions lower than 1 and higher than 0, which is always lower than either of those fractions, one of them being $Pr(\varphi)$.

[66] This is also a countermodel to the conjunction of the claim that justification is the property of being a general consequence of E (as defined by DW), and the claim that general consequence is weakly transitive (DW, p.756). The model in question can be easily modified so that φ_1 and φ_2 are general consequences of E (since $Pr(\varphi_1) = Pr(\varphi_2) > t$), $\varphi_1 \wedge \varphi_2$ is a general consequence of $E \cup \{\varphi_1, \varphi_2\}$ (since $Pr(\varphi_1 \wedge \varphi_2 | E \cup \{\varphi_1, \varphi_2\}) > t$), but $\varphi_1 \wedge \varphi_2$ is not a general consequence of E (since $Pr(\varphi_1 \wedge \varphi_2) < t$).

this type that can be constructed, varying in the cardinality of the set of propositions and in the likelihood of each member of such set.[67]

An important fact about the previous type of model is that any typical example of a lottery can be viewed as a token of such type, since any two members of the set of competing lottery propositions of any such case satisfy (i) and (ii). They satisfy (i) because the probability of any lottery proposition is extremely high (therefore, higher than 0) but below 1. And they satisfy (ii), because the probability of one lottery proposition given another one is below the probability of each of those two lottery propositions, which is below 1, so such conditional probability is also below 1.

Consider, as an example, a fair lottery of 10 tickets with one winner. $\Gamma = \{\varphi_i | \varphi_i = \langle \text{Ticket i loses} \rangle$, for any integer i such that $1 \leq i \leq 10\}$. The probability of a ticket losing is 9/10 (which is higher than 0 and lower than 1), and the probability of any ticket losing given that a different ticket loses is 8/9 (so the probability of any ticket losing is also below 1). Thus, any two lottery propositions in Γ satisfy (i) and (ii). Given this, a countermodel to the aggregativity of sufficient probability can be easily obtained in this case. Assuming that t = .5, we have that:

$Pr(\varphi_1) = .9 > t$, $Pr(\varphi_2) = .9 > t$,..., $Pr(\varphi_{10}) = .9 > t$

[67] The above model, as well as the solution to the paradox offered in section 4, are both inspired on an argument against *Aggregativity* by Richard Foley, which is briefly presented in the next citation:

> Intuitively, it would seem that a person might have very good but not perfect evidence for his belief *p* and have equally good evidence for his belief *q* and yet not have as good evidence for their conjunction. After all, if these beliefs are independent (in the sense that knowing the truth or falsity of one would not help one know the truth or falsity of the other) and if the evidence is such that a person exposes himself to some risk of error in having each belief, then in believing the conjunction the person would at least seem to expose himself to an even greater risk of error. And, it is hard to understand why the likelihood of error might not increase to the point where there is not evidence sufficient to justify a belief in the conjunction. (Foley, 1979)

I find Foley's argument persuasive. And I find rather disappointing that DW largely ignored it. Their only comment about the possibility of rejecting *Aggregativity* is that such a move is "almost generally" considered "too drastic" for solving the lottery paradox (p.758), without pointing out where Foley's argument goes wrong and without trying to explain away the intuitions behind it.

$Pr(\varphi_1 \wedge \varphi_2) = .8 > t$

$Pr(\varphi_1 \wedge \varphi_2 \wedge \varphi_3) = .7 > t$

$Pr(\varphi_1 \wedge \varphi_2 \wedge \varphi_3 \wedge \varphi_4) = .6 > t$

$Pr(\varphi_1 \wedge \varphi_2 \wedge \varphi_3 \wedge \varphi_4 \wedge \varphi_5) = .5 = t$

The probability of each conjunction in the countermodel is calculated by using the general conjunction rule and the relevant conditional probabilities, which are identified by the next line of reasoning: "For each lottery ticket, the probability that it is a loser is 9/10 because there are ten tickets and only one of them can be a winner. But if one of the ten tickets is known to lose, the probability of being a loser for each one of the remaining tickets is 8/9, since there are nine remaining tickets and only one of them can be a winner. And if two of the ten tickets are known to lose, the probability of being a loser for each one of the remaining tickets is 7/8, since there are eight remaining tickets and only one of them can be a winner. And so on." Here, saying "If j tickets are known to lose, the probability of the remaining tickets is..." is equivalent to talking about the probability that a ticket loses given that j tickets lose. This line of reasoning shows, more generally, that if j is the number of tickets known to lose, m is the number of possible losers, and n is the total number of tickets, then the probability of a ticket losing given that j tickets lose is $m\text{-}j\ /n\text{-}j$.[68]

[68] Alternatively, the probabilities of the relevant conjunctions are obtained by means of the following reasoning: "If I have one ticket, there are 1/10 chances that it is a winner and 9/10 chances that it is a loser, since there are nine possible outcomes of the lottery in which my ticket is a loser and only one in which it is a winner. If I had two tickets, the chances that one of those two tickets is a winner would be 2/10, and the chances that I lose would be 8/10, since there are eight possible outcomes of the lottery in which both tickets are losers and only two in which one of them is a winner... If I had five tickets, the chances that I win would be 5/10, since there are five possible outcomes of the lottery in which all of my five tickets are losers and five in which one of them is a winner..." Those probabilities can in turned be used to calculate the relevant conditional probabilities, by means of the definition of conditional probability.

In the type of models considered, each lottery proposition is sufficiently probable (that is, it has a probability above *t*), but the more competing lottery propositions are conjoined, the less probable the resulting conjunction is, up to a point to which the conjunction in question is not sufficiently probable (that is, it has a probability equal to or below *t*).

If the present argument is correct, it is reasonable to think that the proof is not a reductio of the assumption that D is structural, since that would require that every other assumption of the proof is correct, including Aggregativity. Given the argument developed in this section, it is reasonable to think Aggregativity is not correct, so D is no longer the most plausible suspect of generating the problematic dilemma presented by DW.

3. Objections to the Argument Against the Aggregativity of Justification

Let us consider some possible objections to the argument presented in the last section. One first objection starts by noticing the structural parallel between the argument for (III) (i.e., if justification is aggregative, so is support) and the argument for (IV) (i.e., if support is aggregative, so is sufficient probability). Afterwards, it identifies a (presumably) fallacious step in the structure of both arguments. Recall that in the argument for (III), justification is assumed to be aggregative, φ and ψ are assumed to be supported, and $\varphi\wedge\psi$ is assumed to not be supported. Then, the argument makes the crucial stipulation that φ and ψ satisfy every other condition to be justified, so they are justified. Similarly, in the argument for (IV), support is assumed to be aggregative, j and y are assumed to be sufficiently probable, and $\varphi\wedge\psi$ is assumed to be not sufficiently probable. Then, the argument makes the crucial stipulation that φ and ψ satisfy every other condition to be supported. According to the objection, the crucial stipulations in both arguments assume that the next schema is valid (where "C" and "C'" are schematic letters which stand for any monadic predicates applying to propositions):

Schema

Assuming that both φ and ψ have C, which is entailed by C', and that $\varphi \wedge \psi$ does not have C, we can consistently stipulate that each φ and ψ satisfy every other condition required by C' (so φ and ψ have C' too).

As it can be appreciated, C = support and C' = justification in the argument for (III), while C = sufficient probability and C' = support in the argument for (IV). However, the objection continues, *Schema* is not valid, since C' might also entail that, for any x and y, if x has C and y has C, then $x \wedge y$ has C.[69] In the case of the arguments for (III) and (IV), it might be argued that even though support is not aggregative, justification entails an aggregative type of support —let us call it 'support*' — or that even though sufficient probability is not aggregative, having a probability above the threshold required for support* is an aggregative property —let us call such property 'sufficient* probability'.

This objection fails. First, notice that the objection entails that if 'support' in (III) refers to support* (whatever property support* is) and 'sufficient probability' in (IV) refers to sufficient* probability, both (III) and (IV) end up being true (even if my arguments in their favor fail). So, the objection does not really affect (III) and (IV). If the objection is correct, it does, however, affect

[69] The simplest case where this occurs is when C' is defined as follows: for every x, x has C' $=_{def}$ x has C and, for every y, if x has C and y has C, then $x \wedge y$ has C.

(V) – i.e., sufficient probability is not aggregative. This is because if sufficient* probability is stipulated as aggregative, (V) cannot be true. Let us leave aside the discussion of this point after the remaining problems of the present objection are examined.

Second, the claim that justification entails support* rather than support, and the claim that support* entails sufficient* probability, are both controversial claims that need to be argued for, rather than intuitive uncontroversial principles concerning the notions they are about. However, no argument has been presented in their favor. Moreover, even if *Schema* is not generally a valid one, no reason has been provided to think that *Schema* fails in the instances under examination, concerning justification, support, and sufficient probability.

Third, even if those claims are defended, that is not enough to counter my argument, since for the proof to succeed *as a reductio of D being structural*, it should not be based on controversial assumptions about justification, support or the relation between those notions and probability, since in that case the proof might as well be considered to be a reductio of such controversial assumptions rather than the assumption that D is structural. If one intends to show that a reasoning qualifies as *proving* something, one ought to make sure that the premises of that reasoning are not based on controversial assumptions.

Let us consider now the objections that might be leveled against (V). First, let us return to the suggestion that support entails sufficient* probability, where the latter notion is stipulated as aggregative. Notice that, unfortunately, the stipulation entails that only a probability of 1 counts as sufficiently high for support*, since admitting a lower probability threshold for support* would allow the conjunction of all the members of a sufficiently large set of lottery propositions to have a probability below such threshold, even if all its members have a probability above it.

Even if the claim that a probability of 1 is required for justification were somehow defensible, such claim is too controversial to use in the defense of one of the proof's assumptions, since the proof is supposed to be based only on (mostly) uncontroversial assumptions.[70] Moreover, in the present context it is illegitimate to assume that a probability of 1 is required for justification, since in such context DW need to accept just the opposite as correct for their reductio of D being structural to work, as described in section 1.[71] So, whether DW or other philosophers accept that justification requires a probability of 1 in their preferred account of the matter, such views cannot play any part in the defense of the proof's assumptions.

Furthermore, assuming that a probability of 1 is necessary for justification biases the whole discussion against the philosophers proposing formal accounts to solve the lottery paradox. To see this clearly, suppose there is a philosopher that identifies justification with sufficiently high probability given the evidence, where that probability is taken to be lower than 1. Such a view entails that justification is not aggregative, since high probability below 1 is not aggregative. Thus, either the proof fails to apply to such view or it discounts it as false from the start, without argument.

[70] From the outset, the claim is not only controversial, but also highly counterintuitive, since it seems to entail that no proposition can be justified on the basis of a non-deductive inference from the evidence. So, every scientific theory would be counted by that claim as not being justified no matter how much evidence in its favor was available if such evidence did not entail the theory in question.

[71] To be precise, the claim that DW need to assume as absurd for the proof to work as a reductio of D being structural is the claim that Restricted Sufficiency amounts to the claim that having a probability of 1 is sufficient for being justified. However, whatever restriction not having D imposes on a proposition, if having a probability of 1 is a necessary condition for justification, no correct specification of Restricted Sufficiency would admit a proposition with a probability below 1 satisfying its antecedent, in which case Restricted Sufficiency would only admit the trivial reading DW assumed as absurd.

A final objection argues that even if the main argument in section 2 is correct, it is not sufficient to make plausible that the proof is not a reductio of D being structural, since DW presented an alternative reductio using the next restricted version of Aggregativity:

Aggregativity*) If each φ and ψ are justified and $\{\varphi,\psi\}$ is consistent, then $\varphi \wedge \psi$ is justified

Instead of assuming that justification is aggregative, the proof now assumes that it is C-aggregative, where a property is C-aggregative iff a conjunction has it if each of its conjuncts are mutually consistent and they both have it. Then, the proof proceeds as before. Provided that P is structural and sufficient for Q and that Q is C-aggregative, the proof shows that if a proposition φ such that $Pr(\varphi) < 1$ has P, then some proposition ψ such that $Pr(\psi) \leq 1/|W|$ has Q (where W is a finite and very large class of equiprobable worlds). [72] The proof is now perceived as showing that, assuming *D* is structural, a dilemma emerges: either Restricted Sufficiency amounts to the trivial claim that propositions with a probability of 1 are justified, or some proposition with an extremely low probability (that is, $1/|W|$) is justified. Again, both horns of the dilemma are considered as absurd, thus concluding that D is not structural.

This version of the proof is consistent with rejecting Aggregativity. However, (Aggregativity*) can also be shown to fail by a small variation of my argument, substituting aggregativity with C-aggregativity. For this purpose, it is enough to show that sufficient probability is not C-aggregative, which can be done using the same countermodels we already considered. For example, assuming $t = .5$, the conjunction of five lottery propositions (given a lottery of ten tickets) is not inconsistent but is not sufficiently probable even if each of its conjuncts is sufficiently probable. The assumption that $t = .5$ is not necessary for the countermodel to work, since the

[72] For the sub-proof of this result, see DW (2006, pp.771-772).

alternative assumption that *t* ³ .5 suffices. This latter assumption can be intuitively motivated as follows: If the probability of a j given E is not above .5, then j is at most as reasonable as ¬φ given E, which makes it unjustified to accept φ on the basis of E.

Even if the assumption that $t \geq .5$ is rejected (for reasons yet to be provided), a countermodel can be constructed insofar as there is a probability threshold above 0 that is so low that it is necessary to have a probability above it for the relevant proposition to be justified. That much is already accepted by the version of DW's attempted reductio relying on (Aggregativity*), since it assumes as absurd to consider that a proposition with a probability of $1/|W|$ (an "extremely low" probability) can be justified. Given that assumption, we can safely accept that $t \geq 1/|W|$ and construct a corresponding model.[73]

The argument defended in these two last sections provides a basis for rejecting Aggregativity or any problematic restriction of Aggregativity. However, given the intuitive appeal that Aggregativity has, it might be tempting to reject some principle used for supporting its premises. In the next section I provide an independent argument in favor of resisting that temptation.

---

[73] An apparent problem with this way of proceeding is that which number $1/|W|$ is has not been determinately specified, since which class of worlds is W and its cardinality is not determined. However, insofar as DW are inclined to accept that there are cardinalities of classes of (relevant) worlds such that 1 over such cardinality is too small a probability for being justified, a countermodel can be constructed using the assumption that *t* is equal or above that probability. I do not consider that one must go to such extents to construct a countermodel of the intended type, since the assumption that $t \geq .5$ is reasonable enough, and nothing DW have argued so far provides a basis for rejecting such assumption.

4. Solving the Lottery Paradox in Its Most Basic Form

In this section, I examine a type of faulty logical derivation, in which the lottery paradox is already present. The form of the paradox occurring in this type of derivation is also shown to be a more basic form of the lottery paradox than the one concerning the notion of justification. I argue that a solution to the lottery paradox of justification should reflect the solution to the most basic version of the paradox presented here. I then argue that accepting Aggregativity obstructs the solution to the most basic version of the paradox, while rejecting Aggregativity allows us to solve it easily. This is taken to support further the claim that the proof is a reductio of Aggregativity, rather than a reductio of D being structural.

Let us begin with the following case. Consider a finite class C of objects such that $|C| = n$, and a class P such that $|C \cap P| = m$. Let us assume that every member of C has the same chances of being a member of P. Thus, the probability of a member of C being in P is m/n. Let us stipulate that $1 > m/n > t$. Let $\{x_1, x_2, \ldots x_{m-1}, x_m, x_{m+1}\}$ be a class of m+1 different members of C. Finally, let us stipulate that the previous conditions are part of a finite set K of known facts. Consider now the derivation depicted next (where '$H_K(\varphi)$' abbreviates 'Given known information K, it is highly likely that φ', and 'Px' abbreviates 'x is a member of P'):

Faulty Derivation:

1) $H_K(Px_1)$. $\therefore$ Px_1.

2) $H_K(Px_2)$. $\therefore$ Px_2. $\therefore$ $Px_1 \wedge Px_2$.

3) $H_K(Px_3)$. $\therefore$ Px_3. $\therefore$ $Px_1 \wedge Px_2 \wedge Px_3$.

…

m-1) $H_K(Px_{m-1})$. ∴ Px_{m-1}. ∴ $Px_1 \wedge Px_2, \ldots \wedge Px_{m-1}$.

m) $H_K(Px_m)$. ∴ Px_m. ∴ $Px_1 \wedge Px_2, \ldots \wedge Px_{m-1} \wedge Px_m$.

m+1) $H_K(Px_{m+1})$. ∴ Px_{m+1}. ∴ $Px_1 \wedge Px_2, \ldots \wedge Px_{m-1} \wedge Px_m \wedge Px_{m+1}$.

Given that no meta-notion occurs in Faulty Derivation (it is entirely about the probability of particular members of C being members of P, given the information in K), let us call 'object-level' any derivation like Faulty Derivation containing no meta-notion. Clearly, Faulty Derivation is deeply faulty, since in it we end up deriving something known to be false (that is: that more than *m* objects are members of C∩P) exclusively from know facts.

It is easy to see that the lottery paradox occurs in some form in derivations like Faulty Derivation. To see this, let C be the class of lottery tickets of *l*, let P be the class of losing tickets of *l*. Given that there is only one winning ticket in *l*, $m = n–1$ and $m+1 = n$, we obtain a contradiction when $\langle Px_1 \wedge Px_2, \ldots \wedge Px_{m-1} \wedge Px_m \wedge Px_{m+1}\rangle$ is derived at the *m*+1 step of the derivation. Arguably, this form of the paradox (let us call it, an 'object-level' form) is more basic than the one involving the meta-notion of justification (let us call it a 'meta-level' form), since all the conceptual resources needed for the object-level paradox to arise are needed for the meta-level paradox, but not vice versa. Let us see why.

Let us assume that $H_K(\varphi)$ provides enough support for φ to make φ very reasonable and hence, to derive φ, even if $H_K(\varphi)$ does not justify φ (after all, the type of support required by justification should be very demanding, allowing that reasonable propositions fail to fulfill such requirement). Even though the justificatory status of φ is not at issue here (and even if the same happens with the remaining propositions being derived), Faulty Derivation remains paradoxical.

It is puzzling to say what exactly is wrong with it, since both inferring φ from $H_K(φ)$ as well as conjunction introduction (used in the derivations) are reasonable. Interestingly, no principle of justification was needed to understand Faulty Derivation's paradoxical status.

The inverse conceptual relation is not correct. We cannot understand the meta-level paradox involving the notion of justification without indirectly discussing the status of derivations like Faulty Derivation. The assumption that each lottery proposition is justified (a necessary assumption for the paradox to arise at the meta-level) entails that each one of them is supported by the available evidence. This, in turn, entails that each step in a derivation like Faulty Derivation is legitimate when such derivation is used to represent the relations of support among the relevant lottery propositions and the set of available evidence.

The previous result is to be expected, given that the notion of support by the available evidence is more basic than the notion of justification,[74] according to the partial account of the notion presented in section 2. Since justification is to be understood in terms of the notion of support, the object-level paradox in Faulty Derivation is more fundamental than the meta-level paradox involving principles of justification such as Restricted Sufficiency, Aggregativity and

[74] The idea that the notion of support is more basic than the notion of justification might allude to the widely discredited idea that concepts are susceptible to a "breakdown" or decompositional analysis, and that I am assuming that the notion of justification has one such analysis in terms of the notion of support. This is not what I have in mind. There are two ways in which a concept C can be said to be more basic than a concept C' which are compatible with there not being an analysis of the latter in terms of the former. We can say that C is conceptually more basic than C' iff, it is a conceptual truth that every C' is a C, but the inverse claim is not a conceptual truth. Thus, we can say, for example, that the concept of animal is more basic than the concept of cat (because it is a conceptual truth that all cats are animals), and this might be true even if there is no analysis of the concept of cat. Alternatively, we can say that C is conceptually more basic than C' iff grasping/understanding C is required for grasping/understanding C', but not vice versa. For example, it is arguable that the notion of time is required to understand the (ordinary) relation of cause-effect, even though the latter might be primitive. Saying that two concepts are in one of these relations (sometimes called "conceptual priority") seems to be compatible with rejecting that one of them is analyzed in terms of the other. In the present context, I maintain that the notion of support is conceptually prior to (more basic than) the notion of justification in both of the senses explained here (although, only the second sense is required for my relevant argument). It is important to notice, however, that both notions might come apart. For example, it is widely recognized that it is a conceptual truth that knowledge entails justification, but not vice versa, and this might be compatible with the claim that the notion of justification must be understood in terms of the notion of knowledge (Williamson, 2000).

Non-Zero Probability. Thus, the correct solution to the lottery paradox about justification should reflect the correct explanation of what is at fault in the corresponding class of object-level faulty derivations.

Now, since Faulty Derivation is an object-level derivation, no meta-principle involving the notion of justification is employed in Faulty Derivation. For this reason, no such principle can explain why Faulty Derivation is faulty. Rather, the explanation lies on the fact that Faulty Derivation involves some illicit step. Since only the next two inferential rules are used in Faulty Derivation, one of them must be at fault:

H_K-E: $H_K(\varphi)$. Therefore, φ.

$\wedge$−I: φ, ψ. Therefore, $\varphi \wedge \psi$.

When I say that one of these two rules is at fault, I do not mean to say that one of them is logically incorrect (that is, invalid or insufficiently strong, when used in a direct inference), but that it is possible that a proposition φ is logically derived (in a series of steps) from a set of premises K, and yet φ is not supported by K.

Now, although Faulty Derivation's faulty character is not explained by some faulty principle of justification, the correct principles of justification (whichever they are) should reflect the way in which such faulty character is to be avoided. This is because if such faulty character is explained by Faulty Derivation containing an illicit step in which j is derived, this precludes φ from being supported by K, hence precluding φ from being justified (assuming such derivation is the only way in which φ is derived from K, and that K is the only set of available evidence supporting φ).

Now, DW accept Aggregativity, so they would see nothing wrong with $\wedge$-I and would consider H_K-E instead as the rule at fault. Given that H_K-E parallels Sufficiency, it is plausible that DW's explanation of why H_K-E is faulty parallels their strategy to deal with Sufficiency's failure (consisting in incorporating a condition of being undefeated to Sufficiency, thus obtaining Restricted Sufficiency). The strategy they would follow, then, is to restrict H_K-E by a condition of being undefeated, as follows:

+H_K-E. $H_K(\varphi) \wedge \neg D\varphi$. Therefore, φ.

However, this explanation of why Faulty Derivation is faulty is dubious. First of all, +H_K-E is inapplicable unless it is previously determined what property 'D' stands for, a task that is made more difficult if D is not structural. Secondly, the solution is ineffective, since all we need to do for the problem to resurface is replace each step of the derivation employing 'H_K' with a statistical syllogism[75] of the form "The ratio of members of C that are members of P is m/n (where m/n is very close to 1), x_i is a member of C. Therefore, x_i is a member of P." All of the arguments of this form seem to be logically correct, not involving any fallacy or apparent vice.[76] Moreover, each such argument is guaranteed to be a strong one, given the stipulations that, for every *i* such

[75] See (Salmon, 1963) for a presentation of statistical syllogism as a correct inductive inference form. See Pollock (2010) for a more recent probabilistic version.

[76] It might be argued that each one of the statistical syllogisms are not inductively correct derivations, because if they were, they would allow us to end up contradicting the (known) proposition that there is a winning lottery ticket, which is part of the available evidence. Of course, we know that ends up happening, but recognizing that fact does not say why any of the statistical syllogisms are not inductively correct taken individually, but rather point to the obvious fact that they are not *jointly* acceptable.

Notice that appealing to the defeasibility of inductive arguments does not help here, since the proposition that there is a winner does not diminish the support that each of the statistical syllogisms provide to the lottery proposition that is their conclusion (mainly because the fact that there is a winner is part of the very statistical setup which makes each lottery proposition have their current probability).

Finally, notice that if some fault is found in the statistical syllogisms above, then some similar form of correct inference from the evidence should exist, given that i) lottery propositions are supported by the available evidence, ii) being justified entails being supported, and iii) lottery propositions are (arguably) justified (based on the available evidence).

that $x_i \in C$, $Pr(Px_i) = m/n > t$, and that, for some t such that $.5 \leq t < 1$, a probability above t is enough for a derivation to count as strong.

So, while it makes sense to add a condition of being undefeated to high probability (in order for both conditions to be jointly sufficient for justification), appealing to such extra condition is not helpful in accounting for what is supposed to be wrong with the statistical arguments replacing the derivations in which 'H_K' is employed. It is not helpful because the explanation in question should be given at the level of support, and it would be unreasonable to hold that the propositions derived in those statistical arguments are not supported by the premises (all of which are members of K).

Summarizing, the objection presented so far against DW's view on the proof is as follows. The lottery paradox occurs in its most basic form in object-level derivations (about the probability of objects in a certain class having a certain property, given a certain statistical arrangement), not involving any meta-notion, such as the notion of justification. Thus, faulty principles of justification are not the source of the paradox. Rather, the paradox emerges for the notion of justification because such meta-notion is understood in terms of the relations of support between propositions and E. However, any solution to the paradox at the meta-level depends on the solution offered to it at the object-level. And, the suggestion of introducing a condition of being undefeated by the evidence in those principles says nothing about why the relevant object-level derivations are faulty, so DW's view on the proof is inadequate to solve the lottery paradox in its most basic form. Therefore, DW's view on the proof is a red herring with regards to finding a solution to the lottery paradox, since it changes the focus from the epistemic relations that are at the core of the paradox to the meta-principles about justification which are less fundamental than those relations.

In contrast, we can explain why it is not generally the case that ∧-I provides support from E to propositions derived from its use.[77] First, it is clear that the more times ∧-I is applied to conjoin propositions of the form Px_i, the riskier the inference is, up to a point (which occurs exactly at step m+1 of the sequence) where deriving something inconsistent with what is known is guaranteed. But using ∧-I is too risky to provide support before an inconsistency even arises. To see this, consider again the plausible stipulation that, for some t such that $.5 \leq t < 1$, it is necessary that φ has a probability above t, given E, for the derivation of φ from E to be strong enough for φ to be supported by E. Thus, if a conjunction has a probability of .5, deriving it from E via ∧-I is too risky for φ to be supported by E. This explanation not only constitutes a solution to the lottery paradox about the notion of justification,[78] but it also allows us to resist DW's reductio of the assumption that D is structural.

5. Objections to My Proposal and an Alternative to Aggregativity

Let us now consider some objections to the present proposal. The first objection is that if we reject Aggregativity, we lack a general principle providing sufficient conditions for a conjunction to be supported by E when it is derived, by means of ∧-I, from a set of propositions supported by E. But that principle is important for assessing the justificatory status of conjunctions based on the justificatory status of their conjuncts. Thus, if Aggregativity is to be rejected an alternative should

[77] The following explanation is suggested by the intuitions displayed in Foley's argument against (A), referenced in a note above.
[78] This solution to the paradox is defended by (Kyburg, 1961) and (Foley, 1979), among others.

be provided. However, the principle in question is very easy to find given the previous proposal. Just consider the next restriction of Aggregativity, straightforwardly suggested by the explanation of the previous paragraph. For some t such that $.5 \leq t < 1$, we have that:

A+) If each φ and ψ are justified and $\Pr(\varphi \wedge \psi) > t$, then $\varphi \wedge \psi$ is justified

My reason for accepting (Aggregativity+) is as follows. If having a probability above t is a necessary condition for justification (as previously argued), Aggregativity should be restricted to exclude conjunctions that are too unlikely to be justified even if their conjuncts are justified. (Aggregativity*) is too weak a restriction for that purpose. In contrast, (Aggregativity+) is the weaker restriction of Aggregativity that does the work. Since we lack reasons for thinking that a conjunction of justified propositions can fail to be justified even if its probability is above t, (Aggregativity+) is the restriction that is better supported by the reasons discussed so far. (Aggregativity+), however, does not allow DW's reductio to work, for we cannot use (Aggregativity+) to obtain a dilemma one of whose horns says that a proposition having too low a probability (to be supported by E) is nevertheless justified.

Now, it might be objected that (Aggregativity+) is an arbitrary restriction of Aggregativity. There are two ways in which this objection can be understood. The first one is to argue that accepting (Aggregativity+) rather than a different restriction of Aggregativity, or no restriction at all, is unjustified (and, thus, arbitrary). Thus understood, the objection is incorrect, since we already provided the required justification for (Aggregativity+). As we stipulated, t picks out the probability threshold above of which a proposition can have support by the available evidence, which is in turn a necessary condition for being justified. (Aggregativity+) simply says that a conjunction with a probability above that threshold is justified if their conjuncts are, which is the minimum restriction to Aggregativity that is required to obtain a *correct* principle, avoiding the type of countermodels proposed in section 2.

The second way to understand the objection is this. Any restriction of Aggregativity that is useful to assess the justificatory status of conjunctive propositions, based on the justificatory status of their conjuncts, requires a way of identifying non-arbitrarily in which circumstances a conjunction is supported by E when its conjuncts are also supported by E, and in which circumstances the conjunction in question is not so supported. Since it is plausible that *t* stands for a vague threshold (given the intuitive vagueness of the notion of support), the identification of such circumstances is plausibly vague as well. Thus, drawing a line between conjunctions that are supported by E and those that are not, is arbitrary.

However, the vagueness in question is not so severe as to make (Aggregativity+) useless to make assessments of the justificatory status of conjunctions based on the justificatory status of their conjuncts, as it is shown next. Let us begin by considering the set P of all the propositions of the form Px_i, where $x_i \in C$. Remember that, if φ and ψ are members of Π, $Pr(\varphi|\psi) < Pr(\varphi)$, so their probability is mutually dependent. Thus, where $x_1,..., x_j$ are *j* different members of C, $Pr(Px_1 \wedge ... Px_{j-1} \wedge Px_j)$ is calculated by means of the general conjunction rule as follows:[79]

$Pr(Px_1 \wedge ... Px_{j-1} \wedge Px_j) =$

$Pr(Px_1)Pr(Px_2|Px_1)Pr(Px_3|Px_1 \wedge Px_2)...Pr(Px_j|Px_1 \wedge ... Px_{j-1}) =$

$(m/n)(m\text{-}1/n\text{-}1)(m\text{-}2/n\text{-}2))(m\text{-}3/n\text{-}3)...(m\text{-}(j\text{-}1)/n\text{-}(j\text{-}1))$

[79] All the relevant conditional probabilities in this formula are determined as proposed in section 2.

Now, let *t* be the largest probability such that it is determinately the case that φ is supported by E (in context c) only if Pr(φ|E) > *t* (in context c).[80] Identifying the circumstances the objector demands amounts to identifying the smallest integer *k* such that the probability of a conjunction with *k* different members of Π, that is, $(m/n)(m\text{-}1/n\text{-}1)(m\text{-}2/n\text{-}2))(m\text{-}3/n\text{-}3)...(m\text{-}(k\text{-}1)/n\text{-}(k\text{-}1))$, is equal or less than *t*. Given how *t* and *k* were defined, a conjunction $\langle Px_1 \wedge ... Px_{k-1} \wedge Px_k\rangle$ of *k* different members of Π is not determinately supported by E when derived from $\langle Px_1 \wedge ... Px_{k-1}\rangle$ and Px_k by means of ∧-I. This is because it is not determinately the case that $Pr(Px_1 \wedge ... Px_{k-1} \wedge Px_k)$ is high enough for $\langle Px_1 \wedge ... Px_{k-1} \wedge Px_k\rangle$ to be supported by E, even if, individually, both $\langle Px_1 \wedge ... Px_{k-1}\rangle$ and Px_k are determinately supported by E. In contrast, deriving a conjunction of less than *k* different members of Π, by applying ∧-I to them, is determinately supported by E if each of those members is determinately supported by E (since it is determinately the case that the probability of such conjunction is high enough to be supported by E).

The previous response might be objected as follows. Normal epistemic agents are not in a position to identify which probability is *t* (at least, if *t* is defined as the largest probability such that it is determinately the case that φ is supported by E only if Pr(φ|E) is above it). However, even if that is true, the previous proposal shows how to identify non-arbitrarily whether a given conjunction is supported by E or not, provided that its conjuncts are supported by E, in all those circumstances in which we can identify probabilistic values that are determinately high enough to be supported by E or that are determinately too low to be supported by E. If the present objection is correct, it is unreasonable to demand more of a principle about the evidential support that a

[80] The relativization to a context is meant for making room for the possibility that there are different probabilistic thresholds for support in different contexts.

conjunction has in virtue of the evidential support of its conjuncts, given the limitations in identifying probabilistic thresholds of evidential support recognized by that very same objection.

Recapitulating. If the argument developed in this section is correct, it shows that considering DW's proof as a reductio of D being structural rather than as a reductio of Aggregativity misses the epistemic phenomenon that gives rise to the most basic form of the lottery paradox, since such phenomenon has at its core the notion of support by the evidence rather than the notion of justification.

An important point to be considered is that, since it is plausible that the notion of support can be completely captured in terms of logical relations between the target proposition and the available evidence, it is also plausible to suppose that such notion is structural. This is good news for those interested in proposing formal accounts of the notion of support and of the notion of justification. Of course, if the argument developed here is correct, the lottery paradox is solved by finding an appropriate restriction of Aggregativity, so it is unnecessary to find an appropriate restriction of Sufficiency (that is, an adequate specification of D in Restricted Sufficiency) for the purpose of solving the paradox. However, inasmuch as Restricted Sufficiency is plausible independently of its potential role in solving the lottery paradox, the argument developed here frees us to understand D as structural if we wish to do so.[81]

[81] Some classical solutions to the lottery paradox, such as those presented by BonJour (1985) and Ryan (1996), are cases of the general intuition that there are structural features shared by all and only lottery propositions preventing them from being justified. I find that intuition very plausible, even though it is strictly unnecessary for the purpose of solving the lottery paradox if the main arguments developed in this paper are correct. However, if those arguments are correct, they free us to develop such an intuition without the constraints that a solution to the lottery paradox imposes.

6. Conclusion

In this paper, I provided two arguments against Aggregativity, that is, the principle claiming that justification is aggregative. The first argument shows that Aggregativity is incorrect, as well as any restriction of it that could be used in the dilemma posed by DW. The second argument shows that the solution to the most basic form of the lottery paradox (occurring in object-level logical derivations, attempting to provide support by the evidence to the propositions being derived) is obstructed if Aggregativity is maintained, while it is achieved easily if Aggregativity is rejected.

Both arguments are complementary to one another. The first one provides a direct reason to reject Aggregativity, while the second one suggests that not rejecting Aggregativity prevents solving the paradox at its most fundamental level, missing its very core as a consequence. They provide a solid basis for considering the dilemma posed by DW as showing that justification is not aggregative, because support (the most basic notion the paradox is about) is not aggregative and it is plausibly structural.

References

Bondy, P. (2013). How to Understand and Solve the Lottery Paradox. *Logos & Episteme*, *4*(3), 283–292.

BonJour, L. (1985). *The structure of empirical knowledge*. Harvard University Press.

Carnap, R., & Jeffrey, R. C. (1980). *Studies in inductive logic and probability* (Vol. 2). Univ of California Press.

Douven, I., & Williamson, T. (2006). Generalizing the lottery paradox. *The British Journal for the Philosophy of Science*, *57*(4), 755–779.

Foley, R. (1979). Justified inconsistent beliefs. *American Philosophical Quarterly*, *16*(4), 247–257.

Kyburg, H. E. (1961). *Probability and the logic of rational belief*. Middletown, CT, USA: Wesleyan University Press.

Nelkin, D. K. (2000). The lottery paradox, knowledge, and rationality. *Philosophical Review*, *109*(3).

Pollock, J. L. (2010). Defeasible reasoning and degrees of justification. *Argument and Computation*, *1*(1), 7–22.

Pritchard, D. (2007). Anti-luck epistemology. *Synthese*, *158*(3), 277–297.

Ryan, S. (1996). The epistemic virtues of consistency. *Synthese*, *109*(2), 121–141.

Salmon, W. C. (1963). *Logic* (Vol. 29, Issue 2). Englewood Cliffs, N.J., Prentice-Hall.

Sylvan, K. (2014). *On Divorcing the Rational and the Justified in Epistemology*. https://philarchive.org/archive/SYLODT

Talbott, W. (2016). Bayesian Epistemology. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/win2016/entries/epistemology-bayesian/.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press on Demand.

* Section 4's integration into the dissertation project

Section 4 not only restores the possibility of lottery propositions possessing a justification-defeating structural property, but it also carries out the "drastic" move of arguing against the aggregativity of justification (DW, 2006:758). As it turns out, such move has also been made via arguments different from the one presented here by (at least) Kyburg (1961) himself and Richard Foley (1979). The upshot of Sections 3 and 4 is that if lottery beliefs are not justified, their lack of justification can now be accounted for by their possessing either a structural or non-structural property, and if they are justified, we nonetheless have a more robust solution to the lottery paradox. Of course, it can be both the case that *Aggregativity* and *Restricted Sufficiency* are both incorrect, and the conclusions in both sections allow it.

Sections 1-4 have taken up the task of methodically travelling all the main roads to explaining why we fail to know lottery propositions. In the process, it has closed some of the most traveled ones (sections 1 and 3.4), and reopened others (sections 3.3 and 4).

At this point, it could be thought that a practical attitude towards the epistemic status of lottery propositions would be to assume that they are justified, since an alternative and suitable solution to the lottery paradox has been provided. However, on the one hand, this route leaves us with the task of developing an informative account of epistemic luck to explain why we fail to know lottery propositions which, given the results in Section 1, is more challenging than initially thought. On the other, there is still further motivation to think they could be unjustified: a subset of them, legal lottery propositions –that is, litigated claims supported by bare statistical evidence–, do not satisfy epistemic legal standards that seem to be very similar to the justification standard. Standards of this kind, for example, are the "beyond reasonable doubt" standard or the "clear and convincing evidence" standard. Section 5 focuses more generally on epistemic interpretations of standards of proof that aim to account for why legal lottery propositions do not meet some standards of proof.

5. Externalism about Standards of Proof: Two Case Studies of a Faulty Condition

1. Introduction

This paper is concerned with interpretations of legal standards of proof. Standards of proof are legal standards expressing a threshold in the level of "proof" or support that the evidence (admitted at a trial) is to provide the litigated claim in order to convict or hold someone liable. A common way to interpret levels of evidential proof or support for the litigated claim is in terms of degrees of probability of such claim given the admitted evidence. For example, the "beyond reasonable doubt" standard for criminal conviction establishes a threshold of around 90 to 95% of likelihood given the evidence, while the "preponderance of the evidence" employed in civil courts requires above 50% to assign liability.[82]

Understanding the level of proof or support in terms of degrees of probability generates a deeply puzzling situation: statistical evidence alone (for example, in the form of base rates) may render a litigated claim highly probable, satisfying the threshold of a given standard of proof and thus being enough to establish a verdict against the defendant, even though it seems inappropriate to convict or hold someone liable on the basis of such evidence alone. This puzzle has come to be known as "the puzzle of bare statistical evidence" or "the proof paradox", and it is normally presented in the form of two contrasting situations:

[82] See Simon & Mahan (1971). These standards are mainly found in legal systems such as Anglo-American common law. See McCauliff (1982) for other standards of proof in the U.S.

Statistics - Smith vs. Red Cab:

On a Sunday night, Mrs. Smith's car is forced off the road by a taxi and suffers physical injuries. She sues Red Cab Company. The evidence is as follows: She saw that the car forcing her off the road was a taxi, although she was not able to identify its color. Yet, 7 out of 10 taxis circulating on Sunday night are from Red Cab Company, whereas only 3 out of 10 are from Green Cab Company. [83]

Testimony - Smith vs. Red Cab:

On a Sunday night, Mrs. Smith's car is forced off the road by a taxi and suffers physical injuries. She sues Red Cab Company. The evidence is as follows: an independent eyewitness saw a red taxi causing the accident and has testified accordingly.

In Statistics – Smith vs. Red Cab, the probability that a taxi from Red Cab caused the accident given the statistical evidence is .7, which is above the .5 threshold of the preponderance of the evidence standard, if such standard is understood probabilistically. Now, the reliability of eyewitness identification is considered to be around .7 as well.[84] This figure is also above the relevant threshold. Therefore, both mere statistical evidence and testimonial evidence support that Red Cab is at fault, at least enough to satisfy the preponderance of evidence. However, courts would not rule against Red Cab based on the mere statistics that 7 out of 10 taxis at the time of the incident were from Red Cab.[85] Our common sense also seems to confirm that such statistical data

[83] The situation here described is a stylized version of an actual legal case (see Smith, 1945). See also Sargent (1940) for another legal case from which other stylized versions of the lottery paradox arise.
[84] I am here following Fields (2013: 1799).
[85] In the original case there was statistical evidence alone to support that a bus causing the accident belonged to a given company, and the court ruled in favor of the defendant. Its reason was that "while the defendant had the sole

is insufficient to rule against Red Cab. In contrast, we may take eyewitness identification alone to be sufficient (assuming it has been judged relevant and admissible). If these considerations are correct, evidential proof or support should be better understood in terms other than, or in addition to, how probable the claim is given the evidence.[86]

Although some theorists that interpret standards of proof probabilistically have provided answers to the previous puzzle, many others have opted for developing non-probabilistic accounts of standards of proof that aim to explain why bare statistical evidence is insufficient to satisfy a given standard of proof.[87] The latter accounts normally establish a non-probabilistic condition on a standard of proof that is not satisfied by bare statistical evidence. Depending on the type of account, the relevant condition may be a condition that has been proposed to satisfy other epistemic standards, such as justification or knowledge.[88]

This paper focuses on a subgroup of interpretations of legal standards of proof, here called "externalist accounts" – roughly, accounts that establish a condition that may or may not be met regardless of the quality of the evidence for the litigated claim. For two dissimilar externalist accounts of standards of proof, Duncan Pritchard's modal view (2018, 2022) and Michael Blome-

---

franchise for operating a bus line on Main Street, […] this did not preclude private or chartered buses from using this street." (Smith, 1945: 755)

[86] Other situations are used to illustrate different versions of the proof paradox. For example, the Gatecrasher Paradox (Cohen, 1977) describes another civil trial, and the Prison Yard Case (Nesson, 1979: 1192–1193) describes a criminal trial. There are also further reasons to reject probabilism about standards of proof, such as the widely-discussed puzzle in the literature known as "the conjunction paradox" (see Cohen (1977), ch. 5, and Spotswood (2016).)

[87] Koehler and Shaviro (1990), Laudan (2006), and Ross (2021) are among the defenders of interpreting standards of proof in probabilistic terms. For critical surveys on non-probabilistic accounts, see Redmayne (2008), Gardiner (2018), and Pardo (2019).

[88] For example, a subgroup of these accounts propose that for a litigated proposition *p* to meet a standard of proof given the evidence, the factfinder's belief in *p* is to be justified. Different accounts of justification will yield different specifications of such condition. Views of this kind can be found in Ho (2008: 89-99), Buchak (2014), Smith (2018), and Nelkin (2021). Other accounts present conditions on the epistemic standard of knowledge, such as Thomson's "causal condition" (1986), the "sensitivity condition" (Enoch and Spectre, 2019), and a version of the "relevant alternatives condition" (Gardiner, 2020). A third subgroup establishes conditions in which knowledge plays a central role, such as the condition that the factfinder knows that *p* or that the probability that she knows that *p* is above a probability threshold given the evidence. Views of this kind are found in Blome-Tillmann (2017) and Moss (2018).

Tillmann's knowledge-first probabilistic account (2017), I argue that they face independent problems, and that such problems ultimately arise given their shared feature of being externalist.

In section 2, I introduce the internalism/externalism distinction with respect to legal standards of proof. In section 3, I present and discuss a counterexample to Pritchard's modal account, which renders such account too strong. In section 4, I argue that Blome-Tillmann's knowledge-first probabilistic account is also too strong. In section 5, I explain why both accounts are externalist, and argue that this feature best explains why they are susceptible to the problems I raised. I conclude suggesting a desideratum for any account of standard of proof that aims to avoid the problems faced by the modal and the knowledge-first probabilistic account.

2. Internalism and Externalism about Legal Standards of Proof

Internalism about legal standards of proof can be characterized as follows:

> **Internalism LSP**
>
> An account of a legal standard of proof is internalist if and only if it places a condition on the satisfaction of a standard of proof, such that whether it is met is primarily a function of the *quality* of the admitted evidence (and, more generally, the body of information that the court possesses) for the litigated claim.

In Internalism LSP, I take the "quality of the evidence" to denote a cluster of interrelated and intrinsic features that a piece or set of evidence supporting a claim p can have for p to satisfy a given epistemic standard. To explain this concept, let us begin with a pre-theoretic illustration. Our common sense tells us that the "quality of the evidence" for some claim or proposition p

expresses how *good* such evidence is to suspect that *p*, to reasonably believe *p*, to know *p*, or to hold some other epistemic attitude towards *p*. Our common sense also tells us that how good the evidence is will vary depending not only on the target epistemic attitude, but on features of the evidence itself. Consider the following example:

> Your neighbor claims that his niece saw Mrs. Philips striking a neighborhood dog. A day later, you see surveillance footage from another neighbor's house capturing Mrs. Philips striking a dog. Although your neighbor's hearsay report may be good enough evidence for you to be *prima facie* suspicious of Mrs. Philips, it is not good enough for you to reasonably believe that she struck the dog – as is the surveillance footage.

Our common sense tells us that the neighbor's hearsay report is not good enough to reasonably believe that Mrs. Philips committed the act, since it only weakly supports it and additional evidence is required, such as firsthand witness testimonies or evidence for the reliability of your neighbor and his niece's testimonies. It also tells us that the surveillance footage *is* good enough, since it strongly supports that Mrs. Philips committed the act and it is a solid piece of evidence (assuming it hasn't been tampered with). Now, for your neighbor, his niece's testimony may be good enough evidence to come to know that Mrs. Philips committed the act, given relevant information that he possesses about the reliability of his niece's testimony. Given these considerations, we can see that features of a given piece or set of evidence for *p*, such as how solid it is, how varied, how ample, and how strongly it supports a given claim, affect how *good* such evidence is in order to hold an epistemic attitude towards *p*.

The previous pre-theoretic illustration motivates three general features of evidence that determine its quality, which are a version of those proposed by Susan Haack (2014:14) as features of the quality of legal evidence supporting a litigated claim:

i) Support: how strongly the evidence supports a given claim.

ii) Solidity: How solid or reliable the evidence itself is, or how well-supported is the proposition expressing what the evidence is.

iii) Comprehensiveness. How ample and varied is the relevant evidence.[89]

Epistemic accounts of features (i)-(iii) may vary, and it is not the aim of this paper to provide an account of each of them. Rather, the goal of presenting Internalism LSP is limited to providing a very rough understanding of the relevance of (i)-(iii) in accounts of standards of proof for an internalist. To this aim, it should also be noted that (i)-(iii) are interrelated in at least two ways. First, more comprehensive evidence normally affects how strongly it supports a claim. If you had your neighbor's additional information that his niece's testimony is reliable and you had heard her direct story of how she saw Mrs. Philips striking the dog, your set of evidence would better support your belief that Mrs. Philips is guilty. Second, more solid or well-supported evidence normally affects how strongly it in turn supports a claim. The proposition that a video footage *shows* Mrs. Philips striking a dog is better supported than the proposition that the niece *saw* Mrs. Philips striking a dog (one needs only watch the video instead of meticulously testing the reliability of the niece's testimony). That which supports the former proposition –that is, the fact that in watching the video anyone can see that Mrs. Philips is striking the dog– in this case also supports the hypothesis that Mrs. Philips committed the act and supports it better than the niece's testimony and information about the reliability of her testimony.

As mentioned in the introduction, standards of proof express a threshold in the level of proof or support the evidence should provide the litigated claim so as to convict or hold someone

[89] The comprehensiveness of the evidence, understood only in terms of the amount of evidence supporting a hypothesis (as Haack does), has also been called the "weight of the evidence" for a hypothesis (see Nance, 2021). Here I include the variety of evidential sources as constitutive of the evidence's comprehensiveness.

liable. Given that the comprehensiveness and solidity of evidence affect such proof or support, it is reasonable to expect these features to be considered when providing an interpretation of what such thresholds of support are. Therefore, and insofar as support, solidity, and comprehensiveness are features of the evidence proposed as determining its quality, it is (at least in principle) plausible to think of standards of proof as standards of how high the quality of legal evidence is.

The above considerations concerning the notion of *quality* of the evidence allow to reformulate Internalism LSP more precisely as follows:

Internalism LSP

An account of a legal standard of proof is internalist if and only if it places a condition on the satisfaction of a standard of proof, such that whether it is met is primarily a function of:

i) the evidential support for the litigated claim

ii) the solidity or reliability of the evidence itself (or how well-supported is the proposition expressing what the evidence is), and

iii) the comprehensiveness of the evidence.

As previously mentioned, (i)-(iii) extend more generally to the *body of information* that the court possesses for the litigated claim. Reasonably, such body of information includes the admitted evidence. It may also include background knowledge that allows the court to draw relevant and correct inferences about the target claim or to evaluate the evidence itself (legal knowledge, common-sense knowledge, reasoning skills, among others). Finally, the "primarily" qualification is added since there could be unforeseen factors independent of the quality of the evidence that play a role in whether a standard of proof is met.

Now, let us formulate externalism, or non-internalism, about legal standards of proof.

Externalism LSP

An account of a legal standard of proof is externalist if and only if it places a condition on the satisfaction of a standard of proof, such that whether it is met is *not* primarily a function of:

i) the evidential support for the litigated claim

ii) the solidity or reliability of the evidence itself (or how well-supported is the proposition expressing what the evidence is), and

iii) the comprehensiveness of the evidence.

For some accounts, it is clear whether they are internalist or externalist. For example, it is clear that a probabilistic account of standards of proof is internalist. The evidential support for the litigated claim just consists in the probabilistic likelihood of such claim given the evidence, and the reliability and comprehensiveness of the evidence will depend on the methods and processes of data collection. It is also clear, for example, that J.J. Thomson's (1986) causal account of standards of proof is externalist. This account (roughly) proposes that a standard of proof is met only if the admitted evidence is in an "appropriate causal connection" to the fact represented by the litigated claim. Such causal connection is said to "guarantee" (in some, non-specified epistemic sense) that the litigated claim is true. Yet, it is possible for there to be many cases in which the relevant causal connection does not obtain due to the litigated claim being false, despite the admitted evidence being of sufficiently high quality to satisfy a standard of proof.[90]

Given that it seems plausible to think of standards of proof as standards of how good the evidence supporting a litigated claim is, it is in principle reasonable to expect an account of a standard of proof to incorporate features of the quality of the evidence. Of course, such *prima facie*

[90] For a general overview of problems that Thomson's causal theory faces, see Gardiner (2018).

plausibility of internalism is not enough reason to renounce externalist accounts. In fact, we could think of internalist accounts as focusing on the quality of the evidence to provide some form of guarantee that the litigated claim is true, while externalist accounts propose other forms of guarantee that the litigated claim is true (such as Thomson's causal account).

Now, for some accounts, it is not immediately transparent whether they are internalist or externalist. In some cases, for example, some scrutiny suggests they have internalist features, but further evaluation shows that they are ultimately externalist. As will be apparent in the rest of the paper, Pritchard's modal account and Blome-Tillmann's modal account are some such accounts. For now, let us evaluate each of them independently.

3. Pritchard's Modal View

Duncan Pritchard (2018, 2022) proposes that an evidential standard in criminal and civil cases is met only if, given the evidence, not easily could the verdict of conviction or judgement of liability have been wrong. His proposal can be specified as follows: where *p* is a proposition stating that the defendant is at fault for the alleged act or crime,

> Modal Condition (MC)
>
> *p* meets the standard of proof X only if, given the evidence *e* for *p*, it is not an easy possibility that the defendant is wrongfully convicted or held liable.[91]

In this context, Pritchard uses 'wrongfully' in the alethic sense in which a defendant is convicted or held liable even though she is innocent - that is, even though *p* is false. The above condition

[91] Pritchard phrases PMC as the condition that, given the evidence, it is "not an easy possibility that wrongful conviction or wrongful judgement of liability" occurs. The way in which a conviction or judgement of liability is wrongful that is relevant for Pritchard's discussion is when the defendant is innocent but judged guilty or liable. The defendant being innocent is rephrased above as 'p is false' (that is, it is false that the defendant is at fault).

also involves the concept of an *easy* possibility, which is a possible scenario that is overall similar to the actual one. According to Pritchard, we can identify how dissimilar a scenario is from the actual one by the way in which we "naturally order possible worlds", that is, we naturally consider a situation in a possible world w_1 more similar to the actual one than another in w_2 when the actual world would need to change less in order to be w_1 than to be w_2 (2015:443). In this way, actual scenarios count as scenarios overall similar (due to being exactly similar) to the actual world. For example, if in the actual world Red Cab is not at fault but is held liable, it is an easy possibility that Red Cab is wrongfully held liable. In the later versions of his modal account of legal risk (of wrongful convictions or assignment of liability), Pritchard phrases MC as the condition that it is *not epistemically highly risky* that the defendant is wrongfully convicted or held liable. In this formulation, 'epistemically' appears to capture the 'given the evidence' clause, and 'highly risky' that the defendant is wrongfully convicted or held liable captures the condition that it is an easy possibility that she is wrongfully judged this way.

Despite not providing an explicit definition of an easy possibility 'given the evidence' admitted in court, Pritchard's use denotes a possible scenario overall similar to the actual world, where a feature of this scenario is that the court possesses *the same admitted evidence* it possesses in the actual world. As will be apparent below, the clause 'given the evidence' for the modal view does not denote *relative to* the evidence (or more generally, the body of information) admitted in court. Based on an admitted -yet incomprehensive- set of evidence, a court could convict the defendant insofar as the evidence indicates she could not easily have been innocent. In this scenario, the defendant is guilty relative to the evidence, even though in an overall similar world with the same admitted evidence and the same determination of the court, she *is* innocent. As we will see, the relevant fact for MC is that she is innocent in an overall similar world with the same admitted evidence as in the actual one. We will henceforward understand 'given the evidence' as expressing, according to the modal view, a single epistemic constraint on the domain of relevant possibilities: any possibility under evaluation (in particular, easy possibilities) is a scenario in which the court possesses the same admitted evidence it possesses in the actual world.

MC seems to provide a simple guideline when assessing whether a litigated claim meets an evidential legal standard; we need only ask ourselves how much the world would need to change for *p* to be false (and hence for the defendant to be wrongfully convicted or held liable). In the Statistics- Smith vs. Red Cab scenario, our natural ordering of possible worlds says that very little would need to change in the world for a green taxi, instead of a red one, to be passing by and cause the accident, despite the higher probability that a red one was the cause. Hence, very easily could have Red Cab not been at fault and there is a high risk that Red Cab is wrongfully held liable. Thus, MC entails in this case that 〈Red Cab is at fault〉 does not meet the preponderance of the evidence. In contrast, in the Testimony- Smith vs. Red Cab scenario, our natural order of possibilities says that much more than a green taxi passing by would need to happen for the verdict to be incorrect. Maybe the witness lied when testifying seeing a green taxi for some unknown motive, or was 'red-green' color blind unbeknownst to himself, among other series of changes to the actual situation required for the verdict to be incorrect. Therefore, not easily could have Red Cab been at fault and the risk of it being wrongfully held liable is low. MC yields in this case that 〈Red Cab is at fault〉 meets the preponderance of the evidence.

Pritchard's (2018) modal view has been met with objections that point to a misalignment of the strength of the evidence for the litigated claim and the modal proximity of possibilities in which such claim is false. For example, Gardiner (2020) offers a case in which an individual, Ted, would never plan a shooting but there is convincing evidence to inculpate him for doing so. Such evidence is misleading in the sense that it supports Ted planning the shooting but it is false that he planned it. So, it is an easy possibility that Ted is innocent since he is innocent in the actual and all overall similar worlds. Thus, the modal view implies that MC is not satisfied, and that such evidence is to be disregarded. Yet, however misleading the evidence may be, insofar as it is convincing, it would be difficult to reasonably disregard it.[92]

Pritchard (2022) holds that objections as the previous one fail to distinguish between the *actual epistemic risk* of wrongful conviction or judgement of liability and a *court's risk assessment*. The actual (epistemic) risk of wrongful judgement consists in the actual, objective risk of the court giving a wrong verdict when the defendant is innocent. It is a function of the closeness of a possibility (to the actual world) in which the defendant is wrongfully convicted or held liable (2022:4/16). Thus, a wrongful conviction or judgement of liability is actually epistemically risky *when it is an easy possibility that the defendant is innocent but the court convicts her or holds her liable, given the evidence*. For example, the actual epistemic risk of wrongfully convicting Ted is high since it is an easy possibility that he is innocent, given the evidence. In contrast, a court's risk assessment is *its assessment of the closeness of a possibility in which the defendant is innocent, relative to or based on the evidence (or more generally, the body of information available to it)*. For example, the court reasonably judges that the risk of Ted being innocent (and hence wrongfully convicted) is low, relative to the convincing evidence against Ted. The evidence is convincing because it points to Ted being guilty.

With the above distinction in mind, Pritchard defends the modal view. Even though this view implies that the standard to convict Ted is not met (since, given the evidence, it is an easy possibility that he is innocent, and so the actual epistemic risk of him being wrongfully convicted

[92] For objections that also point to convincing but misleading evidence, see Fratantonio (2021). For other objections to the modal view, see Ebert, Durbach, & Smith (2020).

is high), it does not prescribe that the court disregard the misleading evidence against him. The court should not disregard the evidence because the court has only a limited body of information from which to judge, and it is reasonable for them to convict Ted relative to or based on such information. As in other information-relative contexts, it is epistemically appropriate for the court to judge based on the evidence they have at their disposal. Based on the admitted evidence, the court may convict Ted because it judged that he could not easily have been innocent, even though in an overall similar world (with the same admitted evidence and the same determination of the court) Ted is innocent.

Before proceeding, let us highlight that although Pritchard's modal view accounts for why some of the courts' assessments of risk are reasonable *relative* to their evidence, it is not thereby renouncing its commitment to understanding MC in terms of actual epistemic risk. Rather, such commitment is expressed at different points of the presentation of the modal view. For example, Pritchard specifies that "[…] it can both be true that the **actual level of epistemic risk** is high (such that the Blue Bus company **shouldn't** be found liable) and that the epistemically appropriate judgement of the court should be that the level of epistemic risk is low enough to make a liability verdict reasonable. (2022:13/16) Here, the relevant normative claim is that the defendant *should not* be found liable when it is an easy possibility that he is erroneously held liable (i.e., when the risk is high). Reasonably, that the defendant should (or not) be convicted or found liable is a function of whether the claim against her meets the threshold of proof or support required to convict or assign liability. Hence, the above normative claim expresses a commitment to MC in terms of actual risk.

Pritchard's distinction between actual risk given the evidence and a court's assessment of risk relative to the evidence, his specification that the latter is reasonable (given the limited body of information available in court), and his commitment to understanding MC in terms of actual risk given the evidence, raise a few questions. Firstly, why is MC understood in terms of the modal view's account of actual risk, rather than in terms of risk *relative* to or based on the admitted evidence? If a standard of proof is to be a standard of how much the evidence proves or supports a litigated claim, it would be reasonable to expect MC to be a condition of whether the evidence (and body of information, more generally) suggests or indicates the modal proximity of the defendant being wrongfully judged, rather than a condition about actual modal proximity. The evidence in the school shooting case is said to be convincing enough to convict Ted. If we assume that if such evidence is convincing, its quality is high enough to satisfy the beyond reasonable doubt standard, then why is the fact that Ted is actually innocent that which determines that the beyond reasonable standard is not met?

The question underlying the previous considerations is that of whether and how, for the modal view, the quality of the evidence is relevant to meet a standard of proof. Is there any feature that convincing evidence meeting a standard of proof has (as in the Testimony – Smith vs. Red Cab scenario) that convincing evidence not meeting a standard of proof lacks (as the evidence for Ted's guilt), aside from its extrinsic feature of coinciding with actual risk?

Pritchard suggests an answer to this question by specifying that if the court's evidence (or informational basis, more generally) is "reasonably" and "sufficiently comprehensive and sound", then the court's assessment based on the evidence will tend to track and approximate the actual risk of wrongful conviction or judgement of liability (2022:10/16). If this is correct, then the modal view suggests the following diagnosis of the evidence in cases like the shooting case:

> **Diagnosis of Quality)** When the court has convincing evidence for *p* but it is an easy possibility that *p* is false (and thus, it is highly risky that the defendant is wrongfully judged),

if the evidence does not track the actual risk of wrongful conviction or judgement of liability, it is not reasonably or sufficiently comprehensive and sound.

This diagnosis may explain why the evidence for Ted's alleged guilt, despite being convincing, is ultimately inadequate to convict Ted. Such evidence does not track the actual risk of him being innocent (since it *is* an easy possibility that he is innocent while the evidence suggests otherwise). Given Diagnosis of Quality, therefore, such evidence is not reasonably or sufficiently comprehensive and sound. At face value, this diagnosis seems correct, for if the court's evidence included information that generated reasonable doubt over Ted's alleged guilt, the court would realize that it is an easy possibility that Ted is innocent (since the evidence would track the actual risk of wrongful conviction), and the evidence would not convincingly support that Ted is guilty.

Pritchard's commitment to characterizing MC in terms of actual risk, however, remains problematic, as it predicts that *p* does not meet a standard in at least some cases in which there is excellent, high-quality evidence for *p* and more than sufficient to satisfy a given standard. Consider the following scenario:

Framing Twins

Thelma is accused by her acquaintance, John, of defamation. John's lawyer puts forward surveillance footage of the entrance of the building where Thelma works, capturing someone who looks exactly like Thelma sticking a poster with defamatory comments about John on the main door and lighting a cigarette afterwards. Additionally, a co-worker testifies having said goodbye to Thelma while she smoked -as she usually does after work-, while two other co-workers testify that Thelma has always treated John with contempt. There is no further evidence and it is determined that Thelma is guilty of defamation. Unbeknownst to everyone, Thelma has a long-lost identical twin, Louise. Louise and Thelma were separated from birth,

> and thirty years later Louise learns about her sister, investigates her thoroughly, comes to believe her sister has the life she deserved, and deeply resents her for it. She secretly decides to incriminate Thelma by sticking a poster on the main door defaming John, smoking a cigarette afterwards, and behaving like Thelma when saying goodbye to her colleague.

In this scenario, it is an easy possibility that Thelma is wrongfully held liable (since she is actually innocent). Thus, MC entails that the preponderance of the evidence is not met. Notice, however, that there are varied reliable sources of corroborating evidence for the claim that Thelma defamed John: the surveillance footage and three independent testimonies. We could even add further corroborating sources to increase the quality of the evidence, such as another eyewitness leaving the building and chatting briefly with whom he also thought was Thelma. Reasonably, the evidence for Thelma defaming John is excellent and of more than sufficiently high quality to satisfy the preponderance of the evidence. The evidence is solid, since the surveillance footage has not been tampered with, the eyewitnesses were not lying and they chatted with whom they thought was Thelma under normal environmental and psychological circumstances; the evidence is reasonably comprehensive, since it is varied insofar as it includes video footage and testimonial evidence and there is a good number of corroborating sources; and, finally, the evidence strongly supports that Thelma is guilty. Thus, in scenarios like the above, it is not only reasonable for the court to judge that the standard of proof is met, but reasonable for *anyone* who knows all the relevant facts – including the fact that the litigated claim is false– that such standard is met. Framing Twins, therefore, shows that MC is too strong of a condition to meet a standard of proof.

Notice how the shooting case and Framing Twins are similar. In both cases, there is convincing evidence for the litigated claim but such evidence does not track the high risk of the defendant being wrongfully judged. Unfortunately, a central difference among the two cases is that the shooting case is under described. Although we could interpret the evidence in this case as convincing insofar as it strongly supports that Ted is guilty, it isn't clear whether it has the features of being solid and comprehensive. Nonetheless, the description of the shooting case could be in principle as rich as the description of Framing Twins, revealing how a set of evidence may be convincing due to its excellent quality.

It could, of course, be argued that MC is in fact not too strong because in Framing Twins the preponderance of the evidence is not met after all (contrary to our previous considerations). Just as in the shooting case, the evidence for Thelma's guilt is not sufficiently comprehensive and sound (i.e., solid), which may explain why it is ultimately inadequate to hold Thelma liable. Given Diagnosis of Quality, since such evidence does not track the actual risk, it is not sufficiently comprehensive and sound. If the court's evidence were much more comprehensive as to include information that suggested that Thelma has a twin who incriminated her, the court would realize that it is an easy possibility that Thelma is wrongfully held liable, and the evidence would not convincingly support that she is guilty.

Unfortunately, Framing Twins also constitutes a counterexample to Diagnosis of Quality. The evidence does not track the high risk of Thelma being wrongfully held liable, but it is more than sufficiently comprehensive to meet the preponderance of the evidence. Not only is the evidence ample and varied, but it would be unreasonable that to satisfy (or not) a standard of proof the evidence needed to suggest that some atypical, outlandish fact obtains – in particular, the possibility in which Thelma is innocent due to an elaborate, secret, incriminatory plan carried out by a long-lost twin of whose existence no one is aware. In other words, it would be unreasonable to require that the evidence tracked the high risk of Thelma being wrongfully held liable to satisfy (or not) the preponderance of the evidence, given the highly inaccessible and abnormal nature of the facts that account for her innocence. Unless the modal view placed such requirement, there is no good reason to deny that the evidence for Thelma's guilt is sufficiently comprehensive to satisfy the preponderance of the evidence.

Of course, the modal view does not place such a requirement. It does not require that the evidence track the actual risk of wrongful conviction or assignment of liability to satisfy (or not) any given standard. In Framing Twins, doing so would amount to requiring that *relative to* the evidence (or body of information), it is an easy possibility that Thelma is wrongfully held liable. In other words, the evidence would need to indicate or suggest that it is an easy possibility that Thelma is wrongfully held liable. For the preponderance of the evidence not to be met, MC simply requires that it is an easy possibility that Thelma is wrongfully held liable.

Finally, Framing Twins could also be contested by explaining away our intuition that the preponderance of the evidence is met in this case. In judging Thelma liable, the court violates the norm expressed by MC: "a person ought to be judged liable only if it is an easy possibility that she is wrongfully judged liable". Yet, As Pritchard suggested, the court is not blameworthy for thinking that the standard is met, since it could not have done otherwise "given the limited body of information" available to it. Maybe it judged Thelma liable because, in trying to comply with the norm expressed in MC, it determined that it was a far-off possibility that Thelma was innocent. And, contrary to what we think, our intuition that the standard is met is not tracking the propriety of conforming to a norm such as "a person should be judged liable when the evidence is of sufficient quality to hold her liable". Rather, it is tracking the court's blamelessness in judging Thelma liable.[93]

[93] Particular thanks to Michael Blome-Tillmann for discussion on this response.

The above response is reminiscent of a common defense of the knowledge-norm of assertion (roughly, one ought to assert *p* only if one knows *p*).[94] Our intuition that asserting *p* is appropriate in cases of non-knowledge –e.g., when we justifiably believe *p*– is explained away by saying that such intuition does not in fact track the propriety of complying with the norm that one ought to assert *p* only if one justifiably believes *p*, but rather tracks the blamelessness in violating the knowledge-norm of assertion.[95]

Unfortunately, the modal theorist's response is illegitimate insofar as this discussion is disanalogous to the knowledge-norm discussion in two important ways. Firstly, our claim that the preponderance of the evidence is met in Framing Twins is not merely intuitive, as is intuitive that we should assert *p* in some non-knowledge cases. We have a reasonable explanation of why such standard is satisfied. The evidence for Thelma's guilt is of excellent quality –it highly supports that she is guilty, it is reliable, and it is very comprehensive–, and clearly of much better quality than the single testimony in Testimony – Smith vs. Red Cab. If a standard of proof is a measure of how much the evidence proves or supports *p*, and in the Testimony case the evidence is sufficient to hold Red Cab liable, then in cases like Framing Twins the evidence should be more than sufficient.

Secondly, while there is a variety of motivations for the knowledge-norm of assertion,[96] there is only one (epistemic) motivation for MC, and there is reason to suspect such motivation is misguided. The motivation for the modal view consists in its success in accounting for why mere statistical evidence is insufficient to satisfy a standard of proof. In section "5.5.1. Diagnosing the

[94] See Williamson (2000) and McGlynn (2014).
[95] See, for example, Williamson (2000:256) and Hawthorne & Stanley (2008).
[96] Among other motivations, knowledge-first theorists believe that intuitively, we find it inappropriate to assert "*p* and I don't know that *p*", or to assert lottery propositions based merely on probabilistic grounds. It also seems appropriate to ask a speaker that asserts *p* "how do you know that *p*?".

Modal Account", I will argue that such motivation (at least equally or) better supports an internalist version of MC.

For now, let us recapitulate. Framing Twins is a scenario in which the preponderance of the evidence is met, contrary to what MC entails. It is also a scenario in which the evidence is more than sufficiently comprehensive to meet such standard, even though it does not track the risk of wrongful assignment of liability. Since there is no good reason to think the evidence is not sufficiently comprehensive, there is no other feature of the quality of the evidence that the modal view would appeal to in explaining why the standard of proof is not met. In fact, even if the evidence in Framing Twins were not reasonably comprehensive, the modal view does not require that it be reasonably comprehensive in order to track the actual risk of wrongful judgement. Neither does the modal view state that to track actual risk it is sufficient that the evidence be of great quality (including it being reasonably comprehensive). Pritchard only specifies that if the evidence is sufficiently comprehensive, it *tends* to align with the actual risk of wrongful conviction or judgement of liability. This allows there to be cases in which the evidence is of excellent quality but does not track the actual risk of wrongful judgement.

4. Blome-Tillmann's K-First Probabilistic Account

According to Blome-Tillmann's Knowledge-First approach to legal evidence, the standard of proof in civil cases is met when (and only when) given the evidence, the probability that the court knows that the defendant is guilty is greater than 0.5:

Knowledge Rule of the Preponderance of Evidence (KPE)

> *p* meets the standard of proof ('SP' to abbreviate) of the preponderance of evidence iff $P(K_p|e) > .5$. (2017:285)

KPE explains why in Statistics - Smith vs. Red Cab the evidence that it is .7 likely that a red taxi caused the accident does not meet the preponderance of the evidence. According to Blome-Tillmann, in this case the court knows (or is in a position to know), given the evidence, that it does not know that a red taxi caused the accident, which entails that the probability that the court knows that Red Cab is at fault is zero. The explanation for why the court knows that it does not know that Red Cab is at fault is as follows:

> In Smith vs. Red Cab, the court's belief that a red taxi caused the accident is, "if true, true as a matter of luck". For all the court knows, the accident could have easily been caused by a green taxi that belonged to the 30% of taxis circulating in town that night "and there is no evidence whatsoever to eliminate that hypothesis."(2017:286)

The above explains both why the court does not know that a red taxi caused the accident and why the court *knows* that it does not know that a red taxi caused the accident. The court does not know that a red taxi caused the accident because its belief is true as a matter of luck. Of course, such datum by itself does not explain why the court *knows* that it does not know, since one does not always know whether one knows (or does not know) a given proposition. Reasonably, the court *realizes* that, given the evidence, a green taxi could have easily caused the accident (so it would be a matter of luck for a red one to be at fault), and tacitly grasps that this fact prevents it from knowing that a red taxi was at fault. This allows the court to be in a position to know that it does not know that a red taxi caused the accident given the evidence, so that the probability of its knowing such proposition is zero. Thus, the belief that a red taxi caused the accident does not meet KPE.

Despite being ingenious and simple, KPE is not without problems. Let us begin by noticing that for any case where the target legal claim expresses some proposition *p*, either the court A) knows that they know *p*, B) knows that they don't know *p* or C) doesn't know whether it knows or doesn't know *p*. If (A), $P(K_p|e) = 1 > .5$, so *p* obviously meets the preponderance of the evidence standard. If (B), according to Blome-Tillmann's explanation, $P(K_p|e) = 0 < .5$, so *p* does not meet the standard. Both (A) and (B) seem to be unproblematic. However, as I will argue next, KPE is too strong with respect to cases in which we do not know whether we know *p*. Let us see why.

We can safely assume for the present purposes that the court knows (or is in a position to know) that it justifiably believes *p* (abbreviated as 'J_Bp'). In the rest of this section, we will then assume that $P(J_Bp|e) = 1$. We will also assume that knowledge is extensionally equivalent to non-luckily true belief – an assumption with which proponents of K-first need not disagree. One of the strongest motivations for K-first approaches is the long history of failed attempts to provide a non-circular condition which, together with having a true justified belief, would be (necessarily and *a priori*) sufficient to exclude all luckily true beliefs —i.e., the true justified beliefs that, as in Gettier cases, are nevertheless not knowledge— without excluding any case of knowledge. This motivation entails that giving sufficient conditions for knowledge equates with excluding all cases of luckily true beliefs. Inasmuch as it is obvious that knowing entails not having a luckily true belief, the previous motivation entails that the extension of knowledge is equivalent to the class of non-luckily true justified beliefs. This motivation, nonetheless, is consistent with knowledge being a primitive notion as it only imposes an *extensional* desideratum for an analysis of knowledge, while maintaining that such desideratum cannot be met while meeting the *intensional* desideratum of being non-circular (i.e., not relying on the notion of knowledge). Thus, proponents of K-first need not disagree with the claim that knowledge is equivalent with non-luckily true belief as long as such equivalence is not presented as an analysis of knowledge. Moreover, they must not disagree with that equivalence if they accept the previous motivation.

Given the above assumptions, the probability of the court knowing that *p* will be equal to the probability of the conjunction of *p* and the proposition, let us call it '*q*', that there is no condition different from *p* not being true that precludes the belief in *p* from being knowledge. A “gettiered” or “epistemically luckily true” belief, i.e., a belief being true by sheer accident or luck, is the archetype of such condition.[97] We can thus think of *q* as stating the absence of any condition that makes a true belief be epistemically lucky, guaranteeing that the belief in *p* is non-accidentally true.[98] Thus, $P(K_p|e) = P(p \& q|e)$.

Now, let us consider two theoretical possibilities. Either *p* and *q* are independent from each other, or they are in a dependence relation. Both possibilities are problematic. First, let us assume that they are independent from each other, and so, by the multiplication rule for independent events, that $P((p \& q)|e) = P(p|e) \times P(q|e)$. Since the available evidence is gathered only with the purpose of assessing *p*, there will inevitably be cases in which there is no evidence available to assess *q*. Assessing *q* requires evaluating a large range of possibilities under which the court has a luckily true belief in *p* but that are normally considered irrelevant to evaluate *p* itself, e.g., possibilities in which there are unknown environmental conditions disrupting a witness's

[97] In what follows I will use ‘epistemic luck scenarios’ interchangeably with ‘Gettier scenarios.’ There is no theoretically agreed-upon characterization of epistemic luck. Rather, the term is used to describe a range of possible situations in which an epistemic agent has a true (and sometimes justified) belief in a proposition *p* as a matter of sheer coincidence or accident, preventing the agent from knowing *p*. For an overview of characterizations of epistemic luck, see Engel Jr. (2024).

[98] Strictly speaking, *q* states the absence of *any* knowledge-precluding condition different from *p* being false. Gettiered or epistemically luckily true beliefs might turn out to be only a proper subset of the conditions referred to in *q*. I will ignore this specification in what follows, since its truth would bear no negative impact on my arguments and, as will be apparent, would only further strengthen them.

perceptions, possibilities of tampered evidence that nevertheless points to *p* by a stroke of luck, among others. Depending on the situation, the gathered evidence might not say much about whether those possibilities obtain or not, and even if the evidence says something of significance, $P(q|e)$ might still not be high enough so that $P(p|e) \times P(q|e) > .5$. For example, consider the possibility that $P(p|e) = P(q|e) = .7$, in which case $P(p|e) \times P(q|e) = .49$. In many cases, however, the available evidence does not address the conditions under which *q* obtains, since it is only gathered to assess whether *p*. In at least some such cases, $P(q|e) = .5$ due to the limited available evidence, and therefore $P(K_p|e) = P(p \,\&\, q|e) < .5$. To illustrate a case where the evidence does not address the conditions under which *q* obtains, consider first the following scenario:

Misleading Paint

> Raul is accused of vandalizing his neighbor's porch with red paint. He is found liable for property damage based on the following evidence: the plaintiff testified that he and Raul had a heavy disagreement the day prior to the act, and the police found a used can of red paint in a drawer of a cabinet in Raul's study, whose components match those of the paint on the plaintiff's porch. The court concludes that it is more likely than not that Raul used the can of paint found to vandalized his neighbor's porch. Yet, unbeknownst to the court, Raul bought two cans of the same paint and placed one in each of the two top drawers of his cabinet. His son secretly borrowed the one in the left drawer for crafts, and then placed it back, and it was this can the one the police found. Yet, had they opened the drawer on the right, they would have found the can of paint Raul used to commit the vandalizing act.

In this scenario, the court possesses strong evidence for their belief that Raul committed the act: the components of the paint in Raul's left drawer match those of the paint on the plaintiff's porch, and Raul had a motive to vandalize the porch. For the purpose of illustration, let us stipulate that the probability that Raul committed the act given such evidence is, roughly, .9. Nonetheless, the court's belief that Raul committed the act is luckily true, for it was a mere coincidence that the police found the can of paint used by Raul's son and whose components matched the paint on the porch but not the can that Raul used. Of course, given its current evidence, the court does not find it relevant to ask itself whether the drawer on the right contained a can of paint, either used instead of or in addition to the can of paint found by the police. Such question is relevant to us, since we are determining whether the court *knows* that Raul committed the act or if it is a case of epistemic luck.

In the Misleading Paint scenario, the evidence of the court is completely silent about the conditions under which the court's belief is luckily true. Just as easily as the police opened the left drawer, they could have opened the right one and found the target paint, and the evidence does not suggest otherwise, nor does it suggest that there is (or not) a can of paint in the right drawer that Raul might have used to vandalize the porch. Thus, the court's evidence does not seem to render either possibility less likely. In this scenario, then, the probability that the court has a luckily true belief that Raul vandalized the neighbor's porch, given the evidence, is .5.

Now, recall that *q* establishes the absence of *any* condition making a true belief epistemically lucky. So, determining the probability of *q* requires assessing many other possibilities in which the court may have had the luckily true belief in $p = \langle$Raul vandalized the neighbor's porch$\rangle$. For simplicity, let us suppose that the one described above is the only possible situation, so that in this case, $P(q|e) = P(\text{PAINT}^{C}|e) = .5$, where 'PAINT' stands for the set of possible worlds where the above lucky scenario occurs and the superindex '^{C}' for the complement function. As a result, $P(K_p|e) = P(p|e)\times P(q|e) = .9\times.5 = .45$. Thus, KPE is too strong with respect to cases where the court does not know whether it knows *p*, when *p* and *q* are independent. It has the undesirable consequence that, contrary to a reasonable assessment, in some of these cases the preponderance of the evidence is not satisfied.[99]

Consider next what happens if *p* and *q* are dependent. Arguably, the relevant dependence relation is that in which *q* entails *p* since, intuitively, the absence of any condition that makes a true belief be epistemically lucky guarantees that the belief is non-accidentally true. If this is correct, then $P(p|q) = 1$.

To calculate $P(K_p|e)$, recall that we have assumed that $P(J_B p|e) = 1$, and so, that $P(K_p|e) = P((p \,\&\, q)|e)$. Going forward, let us take $K_x=\{w\text{: S knows } x \text{ in } w\}$ (where '*x*' ranges over propositions and '*w*' over possible worlds), and let us take a proposition to be the set of possible worlds in which it is true. Hence, $P(K_p|e) = P((p\cap q)|e)$, where *q* is the set of possible worlds in which *q* is true. Given that $P(p|q) = 1$, using the definition of conditional probability we can derive $P((p\cap q)|e) = P(q|e)$.[100] Thus, $P(K_p|e) = P(q|e)$.

[99] For simplicity, we have assumed that PAINT is the only possibility in which the court has a luckily true belief that *p*. But identifying more possibilities makes KPE's consequence more apparent. Consider an alternative Gettier scenario, PAINT2, identical to PAINT, except for the fact that Raul's *daughter* secretly borrows the can of paint on the left and later places it back. Or consider PAINT3, identical to PAINT, except for the fact that Raul's *sister* borrows the can of paint on the left and later places it back. As is the case with PAINT, the court's evidence is also completely silent about PAINT2 and PAINT3, so the probability of each of those Gettier scenarios given the evidence is .5. Furthermore, PAINT, PAINT2, and PAINT3 are independent from each other: their particular elements of luck may or may not co-occur without affecting the probability of each other's occurrence (imagine the son, daughter, and sister borrowing together the extra can of paint, or only two of them doing so). Due to their mutual independence, by the multiplication rule for independent events, $P((PAINT^C \cap PAINT2^C \cap PAINT3^C)|e) = P(PAINT^C|e)\times P(PAINT2^C|e)\times P(PAINT3^C|e) = .5\times.5\times.5 = .125$. Suppose, for simplicity, that PAINT, PAINT2, and PAINT3 exhaust all possible scenarios in which the court obtains a luckily true belief in *p*. Then, the probability of not being in any lucky scenario (that is, of *q*) would be .125. Therefore, $P(K_p|e) = P(p|e)\times P(q|e) = .9\times.125 = .1125$. While that may be right, there is still strong evidence for *p*, so it does not seem reasonable for the preponderance of the evidence not to be satisfied. Yet, KPE yields that it is not.

[100] The definition of conditional probability is represented by the formula F1: $P(A|B) = P(A\cap B)/P(B)$ or its equivalent F2: $P(A\cap B) = P(A|B)\times P(B)$. To prove that $P(q|e) = P((p\cap q)|e)$, let us first exemplify F2 with *p* and *q* so that $P(p\cap q) = P(p|q)\times P(q)$. Since $P(p|q) = 1$, $P(p\cap q) = 1\times P(q)$, and so $P(p\cap q) = P(q)$. Of course, we can add any further set *x* to both relata, so that $P(p\cap q\cap x) = P(q\cap x)$. With this in mind, notice that $P(p\cap q\cap e) = P(q\cap e)$. If we exemplify F1 with *p*, *q*, and e, so that $P((p\cap q)|e) = P(p\cap q\cap e)/P(e)$, and substitute $P(p\cap q\cap e)$ with $P(q\cap e)$, we can derive $P((p\cap q)|e) = P(q\cap e)/P(e)$. As it turns out, $P(q\cap e)/P(e)$ is nothing else but the probability of *q* given the evidence e. We can see this by exemplifying F1 once more, now with *q* and *e*, so that $P(q|e) = P(q\cap e)/P(e)$. Therefore, $P(q|e) = P((p\cap q)|e)$.

However, KPE is also too strong when *q* entails *p*, as it is plausible that there are at least some cases in which there is good evidence for *p* to intuitively satisfy the preponderance of the evidence and yet $P(q|e) \leq .5$ – that is, cases in which there is *at least one* Gettier scenario whose probability given the evidence is $\geq .5$. Let us begin by stating some assumptions. For any situation in which *q* entails *p* and the court does not know whether they know *p*,

1. $P(K_p|e) = P(q|e)$ *Ass.*

Let us take 'GETT' to designate all Gettier scenarios, and (for simplicity) let us assume that GETT is finite. Consider some partition of GETT, where we designate each of its subsets with 'G_j' and where '$_j$' stands for a natural number, such that $GETT = (G_1 \cup G_2 \cup \ldots \cup G_n)$. We then have that:

2. $P(GETT|e) = P((G_1 \cup G_2 \cup \ldots \cup G_n)|e)$ *Ass.*
3. $P((G_1 \cup G_2 \cup \ldots \cup G_n)|e) = P(G_1|e) + P(G_2|e) + \ldots + P(G_n|e)$ *Addition Rule*

Since we are understanding *q* as the absence of any condition making a true belief epistemically lucky, $q = GETT^C$, where '^{C}' stands for the complement function. With this in mind,

4. $P(q|e) = P(GETT^C|e)$ *Ass.*

The probability of not being in any Gettier scenario is straightforwardly calculated using the complement rule of probability:

5. $P(GETT^C|e) = 1 - P(GETT|e)$ *Complement Rule*
6. $P(GETT^C|e) = 1 - P((G_1 \cup G_2 \cup \ldots \cup G_n)|e)$ *2,5*
7. $P(GETT^C|e) = 1 - (P(G_1|e) + P(G_2|e) + \ldots + P(G_n|e))$ *3,6*

Now, let us take some Gettier scenario g_i (where $1 \leq i \leq n$) in GETT. Note that there will be a partition such that one of its subsets, G_i, has g_i as its sole member, given some defining property of *g* that other Gettier scenarios lack. Problems will arise if there is at least *one* G_i whose probability given the evidence is $\leq .5$.

8. $P(G_i|e) = .5$ (where $1 \leq i \leq n$ and $G_i \subseteq GETT$) *Ass.*
9. $P(GETT^C|e) = 1 - (.5 + P(G_1|e) + P(G_2|e) + \ldots + P(G_n|e))$ *7,8*
10. $P(GETT^C|e) \leq .5$ *9*
11. $P(q|e) \leq .5$ *4, 10*
12. $P(K_p|e) \leq .5$ *1, 11*

We can infer from the above that $P(K_p|e) > .5$ would require the truth of the strong and unjustified claim that *there is no* subset of *any* partition of GETT whose probability is equal to or above .5. But the claim that there is no such subset is not justified. In contrast, it is more reasonable to expect there to be cases in which the available evidence in court is silent about at least one Gettier scenario, g_i, where $G_i = \{g_i\}$ and $P(G_i|e) = .5$. To this point, consider how easy it was to identify scenario PAINT, about which the available evidence in the Misleading Paint case is silent. And while the evidence in these cases may be silent about a given Gettier scenario, they may nonetheless include strong evidence for *p*. However, KPE will render *p* below the threshold of the preponderance of the evidence in any such case.

A proponent of KPE could point out that the above proof wrongly assumes that the probability of at least one Gettier scenario given the evidence is $\leq .5$, that is, it wrongly assumes (8). Gettier scenarios are cases in which something unusual happens that prevent the court from knowing *p*. However, the probability that nothing unusual has happened may be high when accurately measured by the evidence. For example, consider the Testimony – Smith Vs. Red Cab case. According to Blome-Tillmann, "the probability that nothing unusual has happened" is measured by the average reliability of witnesses in court, which makes P(*p*|e) be around .7 (2017:286). Witness testimony seems to render the possibility in which a green taxi caused the accident far off, in contrast with mere statistical evidence, which allows such possibility to be an easy one. And it is precisely due to the testimony's reliability that the world would need to change substantially for the witness' testimony to be false. Thus, the probability that nothing unusual has happened given the witness' testimony is around .7 as well. Consequently, the probability that the court knows *p* given the witness' testimony is around .7.

In response, let us first acknowledge that in the case of Testimony – Smith Vs. Red Cab, witness testimony that a red taxi caused the accident does seem to render the possibility in which a green taxi causes the accident far off, in contrast with the mere datum that 7 out of 10 cars passing by that night were red. Nonetheless, the kinds of unusual events that factor in studies regarding reliability figures are cases in which the testimony is false, involving distortion of reality (resulting from factors such as poor viewing conditions, temporary expectations, or racial and gender bias), compromised memories, testimony resulting from suspect confirmation or suggestive procedures, among others.[101] In such unusual cases, the target claim would be false due to the testimony being false. However, Gettier cases are unusual in a quite different way: they are cases in which the target claim is (luckily) true, and it is natural not to consider them in such studies. Therefore, Gettier cases remain outside the scope of cases in which "something unusual has happened" in the way that unusual events occur where witness testimony fails.

[101] For a discussion on the reliability of eyewitness testimony, see Roberts (2014).

Moreover, KPE would remain problematic even with an assumption weaker than (8). If there is *some* partition of GETT such that the probability of each of its members is below .5 and above 0, and the sum of their probabilities is equal to or higher than .5, the result of subtracting this sum from 1 will be equal to or lower than .5. In such cases, $P(GETT^C|e) \leq .5$, and since $P(GETT^C|e) = P(q|e) = P(K_p|e)$, then $P(K_p|e) \leq .5$.

Another line of defense of KPE is to explain away our intuition that the evidence in Misleading paint is sufficient to find Raul liable for property damage. In judging Raul liable, the court violates the norm expressed by KPE: "a person ought to be judged liable iff $P(K_p|e) > .5$". However, the court is not blameworthy for violating such norm since it could have not done otherwise. Maybe it judged Raul liable because, in trying to comply with the norm expressed in KPE, it determined that something very unusual would need to have happened for it *not* to know that Raul committed the act, given the evidence. Contrary to what we think, our intuition that the preponderance of the evidence is met is not tracking the propriety of conforming to a norm such as "a person should be judged liable when there is good evidence to support that she committed the act". Rather, it is tracking the court's blamelessness in judging Raul liable.

The previous defense mirrors the defense of the knowledge-norm of assertion against the intuition that asserting *p* is appropriate in cases of non-knowledge.[102] And, as has been done with respect to the latter defense, we can evaluate the excuse that renders the agent non-blameworthy in Misleading Paint.[103] In this scenario, the court is said to be excused from blame because, given its evidence, it could not have but determined that something very unusual would need to have happened for it not to know that Raul is guilty. However, there is reason to think such excuse is illegitimate.

---

[102] See the end of section 5.3 for a description of such defense.
[103] See Gerken (2011) for an evaluation of excuse responses from defenders of the knowledge-norm of assertion.

Let us first distinguish between stating i) "something unusual would need to happen for not-*p* to be the case" and ii) "something unusual would need to happen for the court not to know that *p*", that is, what happens for the court to be lucky in getting it right about *p*. The two pieces of evidence in Misleading Paint seem to suggest that something unusual would have to have happened for not-*p* to be the case, that is, for Raul to be innocent. This is expected since the evidence supports that Raul is guilty.

However, even though the evidence suggests (i), it is silent about (ii). What makes the court lucky in getting it right –that is, the unusual event that prevents it from having knowledge– is that it was a mere coincidence that the police found the can of paint Raul's son used and that its components matched the paint on the porch but did not find the can that Raul used. The evidence suggests absolutely nothing about there being or not a second can of paint used by Raul's son – let alone that such can was (or wasn't) the one that the police found while the undiscovered can used to vandalize the porch was in the adjacent drawer. So, the evidence does not make it likely or unlikely that such coincidence obtains: given the evidence, it is .5 likely that it obtains.

The consequence above might be difficult to accept when confusing (i) with (ii), since Gettier scenarios are themselves, regardless of the evidence for *p*, unusual and surprising (after all, they are cases in which it is a mere coincidence that our beliefs are true, even if they are justified). But notice that in many Gettier cases in which a subject S justifiably believes *p*, the evidence that justifies S's belief is silent about the lucky event that prevents S from knowing that *p*. For example, consider the following feature of the evidence in the first case presented by Edmund Gettier (1963). Smith's evidence for thinking that the person who has ten coins in their pocket will get the job is the testimony of the boss that Jones would get the job. Such testimony suggests absolutely nothing about how many coins Smith has in his pocket and the rest of the elements which together prevent

Smith from knowing the target proposition. So, in having evidence to justifiably believe *p*, Smith has reason to think that he knows *p unless* he is in a Gettier scenario. For, as in other Gettier scenarios, luck escapes the evidence. If it didn't, we would be able to avoid them.

Given the previous considerations, insofar as the court in Misleading Paint had good evidence to believe *p*, it had good reason to think that it knew *p unless* it were in a Gettier scenario. The court could not have but determined that given the evidence, something very unusual would need to have happened for Raul to be innocent (that is, (i)). Yet, it could have determined that given the evidence, it equally could or could not be in a Gettier scenario (that is, not-(ii)). Thus, it is not excused from blame.

Most importantly, it is worth highlighting that Misleading Paint is only one of many lucky scenarios whose probability given the evidence would need to be calculated when evaluating whether there is *no* condition that makes the court's true belief be epistemically lucky. The outlook is bleak.

Given the problems for KPE presented above, a probabilistic Knowledge-First analysis of standards of proof is better off at least renouncing its probabilistic element. Insofar as KPE requires the probability that the court knows *p* given the evidence, it requires that the probability of there not being any condition that makes a true belief be epistemically lucky be above .5, given the evidence. Such requirement is too strong to impose on the preponderance of the evidence. As for now, the floor remains open to an alternative, non-probabilistic Knowledge-First analysis.

5. Diagnosing the Modal and the K-First Probabilistic Account

In what follows, I will argue that the modal account and the knowledge-first probabilistic account are externalist, and that being externalist is what best explains why they are too strong.

1. Diagnosing the Modal Account

As seen in section 3, for the modal view it is not necessary that the evidence track the actual risk of wrongful judgement to meet a standard of proof. And it is neither necessary nor sufficient – contrary to first appearances, given Diagnosis of Quality– that the evidence be of good quality in order for it to track the actual risk of wrongful judgement. Insofar as the actual risk condition stands in no logical relation to the evidence being of a certain quality, a standard of proof given MC can be or cannot be met regardless of the quality of the evidence for a given litigated claim. Thus, the modal view's MC account is externalist.

Note once again that, under the modal view, the preponderance of the evidence is met in cases like Testimony – Smith vs. Red Cab but not in cases like Framing Twins, despite the higher quality of the evidence in the latter (in which the evidence is varied and there are more corroborating sources). An explanation of why the standard is met in the former but not in the latter ultimately ends up resting on the fact that the defendant in Framing Twins is actually wrongfully held liable whereas the defendant in Testimony – Smith vs. Red Cab is not. That is, it ends up resting on a fact of how the world is, regardless of how good the evidence for the litigated claim is. Once again, for the modal view, the quality of legal evidence appears to play no role in whether a standard of proof is met.

Unfortunately, the modal view's externalist feature best accounts for why MC is too strong. MC is too strong because it entails that a standard is not satisfied merely because there is a fact that makes it an easy possibility that the litigated claim is false, even when the evidence for the claim is excellent and, reasonably, more than sufficient to convict or hold the defendant liable. In allowing that a fact independent of the evidence determine whether a standard of proof is met, MC opens the door to counterexamples (such as Framing Twins).

Contrast the previous result with a modal view, for example, whose condition to satisfy a given standard of proof were that MC*) the risk of wrongful conviction is low *relative* to the evidence and body of information more generally. Such an account can also explain why statistical evidence is not *good enough* to establish a verdict, while other forms of evidence (such as testimony) can. In Statistics – Smith vs. Red Cab, it is not only an easy possibility that Red Cab is wrongfully held liable given the evidence. It is also an easy possibility *relative* to the evidence and body of information: for all we know, despite being unlikely that Red Cab is innocent, a green taxi could have easily caused the accident, which suggests the need for further information to determine if the preponderance of the evidence is met. Similarly, in the Testimony – Smith vs. Red Cab scenario, the possibility in which Red Cab is innocent is not only a far-off possibility given the evidence, but it is also a far-off possibility *relative* to the evidence or body of information: the court's body of information suggests a lot would need to change for the testimony to be false. MC* provides an response to the proof paradox at least as satisfying as MC, without being subject to counterexamples constituted by scenarios where the evidence plays no role in whether a given proof standard is met.

Notice that MC* is internalist, and it is its internalist features that also account for why we think that statistical evidence is not good enough to convict or hold someone liable. The randomness in statistical evidence such as the fact that 7 out of 10 taxis circulating in town the night of the accident suggests that a green taxi could have very easily been passing by and caused the accident. This reflection is possible by our background knowledge of how random results come about, and our understanding that the evidence in Statistics – Smith vs. Red Cab is a case of randomness. Such considerations also bring to light that such type of statistical evidence is not comprehensive enough, as it very easily allows wrongful judgement of liability.

Interestingly, in both the statistics and testimony scenarios i) there is a high or low risk of wrongful assignment of liability, and ii) the evidence (or body of information) suggests that there is such a high or low risk. Of course, the modal view relies on our intuitions to support (i). In this way, the question of why statistical evidence is not "good enough" to establish a verdict is rather a question about actual risk, disguised as a question about the quality of the evidence. I suggest it is worth exploring, however, whether our intuition that statistical evidence is not good enough to establish a verdict supports (ii) rather than (i): statistical evidence is not good enough when our evidence itself, along with our background information, suggest that very little would need to change for a green taxi to have caused the accident. In this way, the question of why statistical evidence is not good enough to establish a verdict appears to remain a question about the quality of the evidence. If our intuition supports (ii) better or equally than (i), the modal theorist may need to reassess relying on such intuition to support her view.

2. Diagnosing the Knowledge-First Probabilistic Account

At first glance, the Knowledge-First Probabilistic condition expressed in KPE is externalist. To satisfy the preponderance of the evidence standard, the probability of the court knowing that *p* given the evidence should be above .5. Such condition requires evaluating the evidential support

–understood in probabilistic terms– *for the court's knowledge that* p, not the evidential support for p. However, things are not that simple.

If given KPE, whether the preponderance of the evidence is met is a function of the quality of the evidence *for p*, then whether $P(K_p|e) > .5$ depends on i) how much e supports p, ii) how solid or well-supported e is, and iii) how comprehensive e is to support p. It could be argued that in the knowledge-first probabilistic diagnosis of Testimony – Smith vs. Red Cab, the fact that $P(K_p|e) > .5$ depends on (ii). Note first that Blome-Tillmann states that if the court has knowledge that p, p is true and "nothing unusual has happened for the court not to know p" or it is not the case that "if p is true, it is true as a matter of luck". Secondly, according to Blome-Tillmann, the probability that nothing unusual has happened for the court not to know p:⟨Red Cab is at fault⟩ is measured by the reliability figures of eyewitness testimony (2017:286).[104] Such reliability figures are thus taken to constitute the evidence e not only for p, but also for q=⟨nothing unusual has happened for the court not to know that p⟩, where the support e provides q is taken to be $P(q|e)=.7$ (just as is the case that $P(p|e)=.7$). From this datum, Blome-Tillmann infers that $P(K_p|e)=.7$. Thirdly, e expresses a prior probability of reliability of eyewitness identification, $P(e)=.7$, which models how well supported e itself is. Given that $P(e)=.7$, that e supports q such that $P(q|e)=.7$, and that $P(K_p|e)$ depends on $P(q|e)$, then $P(K_p|e)$ depends on how well supported e itself is. In other words, the

[104] It is not clear whether Blome-Tillmann provides an argument to support such claim, and it is also unclear whether the claim is true. It is clear, for example, how the reliability figures of eyewitness identification would constitute evidence for a proposition like r=⟨nothing unusual has happened so that the testimony that p were false⟩. As mentioned in section 5, the kind of unusual events that are normally relevant to these discussions are events concerning the reliability of eyewitness identification (e.g. those that result from factors such as poor viewing conditions, compromised memories, biases, among others). But q makes reference to unusual events concerning the lack of knowledge of the court. The domain of that class of events could be different from that of the class of unusual events expressed in r. For example, it could also include unusual events in which the court does not believe p for unusual circumstances (assuming that believing p is a necessary condition for knowing p).

probability that the court knows that *p* depends partly on (ii). This is an internalist feature of the knowledge-first probabilistic account.

A second internalist feature of the account is that $P(K_p|e)$ is partly a function of (i), insofar as the probability that the court knows *p* given the evidence is equivalent to the probability of the conjunction of *p* and *q*, $P(K_p|e) = P(p \,\&\, q|e)$, regardless of whether *p* and *q* are independent from each other. The fact that whether $P(K_p|e) > .5$ is met depends on (i) and (ii) gives us reason to consider it internalist.

However, there are reasons jointly conclusive to classify such account as externalist. Firstly, whether $P(K_p|e) > .5$ is not a function of (iii). Even though how comprehensive *e* is to support *q* is relevant to satisfy such condition, how comprehensive *e* is to support *p* is of little or no relevance. As explained in section 4, whether *e* supports that there are no Gettier scenarios in place requires in many cases (at least in those in which we do not know whether we know that *p*) evaluating a wide range of scenarios about which *e* is silent. If *e* is silent about such scenarios, it is not comprehensive enough to evaluate them. As a result, the probability of not being in a Gettier scenario ends up being significantly low – or at least low enough to easily render the probability of one knowing that *p* below .5. In other words, the fact that *e* is not comprehensive enough to support *q* is either the predominant or the only factor in explaining why in many cases $P(K_p|e) \leq .5$. It is the predominant factor when *p* and *q* are independent from each other, for we may have situations in which, for example, $P(K_p|e) = P(p|e) \times P(q|e) = .9 \times .5 = .45$. This makes how comprehensive *e* is to support *p* of little relevance for $P(K_p|e)$, since *e*'s support of *p* is itself of little relevance for $P(K_p|e)$. And, when *q* entails *p* (that is, when $P(p|q) = 1$), the fact that *e* is not comprehensive enough to support *q* is the only factor, for we may have situations in which $P(K_p|e)$

$= P(q|e) \leq .5$. In these cases, the comprehensiveness of *e* to support *p* is irrelevant for $P(K_p|e)$, since *e*'s support of *p* is itself irrelevant for $P(K_p|e)$.

Secondly, the above considerations also bring to light that *e*'s support of *p* is of little or no relevance to whether $P(K_p|e) > .5$ (in cases in which we do not know whether we know *p*). Thus, even though KPE expresses a condition whose satisfaction sometimes depends partly on (i) –that is, when $P(K_p|e) = P(q|e)$–, at other times it does not depend on (i) –that is, when $P(K_p|e) = P(q|e)$. Similarly, the above considerations also bring to light that *e*'s support of *q* is either the predominant or the only factor in explaining whether $P(K_p|e) > .5$. The probability of *p* given *e* may be very high, but the probability of *q* given *e* in cases where *e* is silent about *q* significantly decreases $P(K_p|e)$ to a number below the threshold to meet a standard of proof. Ultimately, then, the central factors that determine whether the knowledge-first probabilistic condition is met are external to the quality of the evidence for *p*. Thus, KPE is externalist.

KPE's externalist feature also best explains why this condition is too strong. As we saw in section 5, KPE is too strong because there are cases in which the probability of not being in any Gettier scenario is not above .5, despite the evidence for the litigated claim being excellent to convict or hold the defendant liable. Those scenarios are problematic because the quality of the evidence for *q* is very thin. But they arise in the first place because the focus of KPE is the quality of the evidence for *q*. In both the statistics and testimony versions of Smith vs. Red Cab there seems to be a strong connection between the quality of the evidence for *q* and the quality of evidence for *p*, which renders KPE a promissory account to satisfactorily explain why in the statistics (and not in the testimony) version the evidence for *p* is not good enough for liability. However, the promissory aspect dissipates once discovered that KPE allows for other cases whose evaluation completely, or almost completely, ignores aspects (i) and (iii) of the quality of the evidence for *p*.

6. Concluding Remarks: An Internalist Desideratum

We have seen that by largely ignoring the quality of the evidence for the target litigated claim, two externalist accounts of standards of proof –the modal and the knowledge-first probabilistic accounts– are open to counterexamples in which there is excellent evidence to convict or assign liability while they entail that the proof standard is not satisfied. Whether other externalist accounts suffer the same fate may require identifying them as externalist and arguing for a pattern – for example, that they are all open to counterexamples of the type described above. Interestingly, at least Thomson's causal account is subject to counterexamples in which, as mentioned in section 2, the causal condition is not satisfied even though there is good enough evidence to establish a verdict. In any case, the present evaluations of the modal and knowledge-first probabilistic accounts give rise to a point in question concerning externalism about standards of proof: are externalist accounts, in virtue of being externalist, subject to counterexamples in which they entail that a proof standard is not satisfied but the evidence is of sufficiently high quality to satisfy it?

A working hypothesis answers the previous question affirmatively and suggests an internalist desideratum for an interpretation or account of a standard of proof: *an account of a legal standard of proof should specify conditions whose satisfaction is mainly a function of the quality of the evidence for the litigated claim*. A careful analysis of the virtues and faults of such desideratum will depend on an analysis of the virtues and faults of internalism about legal standards of proof. As for now, there is legitimate motivation to begin evaluating internalist and externalist accounts of standards of proof *qua* internalist and externalist.

References

Blome-Tillmann, M. (2017). 'More Likely Than Not' - Knowledge First and the Role of Statistical Evidence in Courts of Law. In Adam Carter, Emma Gordon & Benjamin Jarvis (eds.), *Knowledge First: Approaches in Epistemology and Mind*. Oxford, UK: Oxford University Press. pp. 278-292.

Buchak, L. (2014). Belief, Credence, and Norms. *Philosophical Studies*, 169(2): 285– 311.

Cohen, L.J. (1977). *The probable and the provable*. Oxford: Clarendon Press.

Ebert, P., Durbach, I., & Smith, M. (2020). Varieties of risk. *Philosophy and Phenomenological Research*, 101, 432–455.

Engel Jr., M. (2024). Epistemic Luck. In The Internet Encyclopedia of Philosophy. https://iep.utm.edu/epi-luck/

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.

Fields, K., L. (2013). Toward a Bayesian Analysis of Recanted Eyewitness Identification Testimony. *New York University Law Review* 88(1769): 1769-1801

Fratantonio, G. (2021). Evidence risk, and proof paradoxes: Pessimism about the epistemic project. *International Journal of Evidence and Proof*, 25, 307–325.

Gardiner, G. (2018). Legal Burdens of Proof and Statistical Evidence. In David Coady & James Chase (eds.), *The Routledge Handbook of Applied Epistemology*. New York: Routledge.

Gardiner, G. (2020). Relevance and risk: How the relevant alternatives framework models the epistemology of risk. *Synthese* 199 (1-2):481-511.

Gerken, M. 2011. Warrant and Action. *Synthese* 178: 529–47.

Haack, S. (2014). *Evidence Matters: Science, Proof, and Truth in Law*. Cambridge University Press.

Hawthorne, J. and Stanley, J. (2008). Knowledge and Action. *Journal of Philosophy* 105: 571–90.

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

Koehler, J. J., & Shaviro, D.N. (1990). Veridical Verdicts: Increasing Verdict Accuracy Through the Use of Overtly Probabilistic Evidence and Methods. *Cornell Law Review*, 75(2): 247–279.

Laudan, L. (2006*). Truth, Error, and Criminal Law: An Essay in Legal Epistemology*. Cambridge: Cambridge University Press.

McCauliff, C.M.A. (1982). Burdens of Proof: Degrees of Belief, Quanta of Evidence, or Constitutional Guarantees?, 35 *Vanderbilt Law Review* 1293.

McGlynn, A. (2014). *Knowledge First?* New York: Palgrave Macmillan.

Moss, S. (2018). *Probabilistic Knowledge*. Oxford: Oxford University Press.

Nance, D. (2021). Weight of Evidence. In *Philosophical Foundations of Evidence Law*. Oxford University Press.

Nelkin, D. (2021). Rational Belief and Statistical Evidence: Blame, Bias and the Law. In *Lotteries, Knowledge, and Rational Belief*. Cambridge University Press.

Nesson, C. (1979) Reasonable doubt and permissive inference—the value of complexity. *Harvard Law Review* 92: 1187–1199.

Pardo, M.S. (2019). The Paradoxes of Legal Proof: A Critical Guide. *Boston University Law Review*, 99.

Pritchard, D. (2015). Risk. *Metaphilosophy* 46 (3):436-461.

Pritchard, D. (2018). Legal risk, legal evidence and the arithmetic of criminal justice. *Jurisprudence* 9 (1):108-119.

Pritchard, D. (2022). In defence of the modal account of legal risk. *Synthese* 200 (4):1-16.

Redmayne, M. (2008). Exploring the proof paradoxes. *Legal Theory* 14 (4):281-309.

Roberts, N. (2014). *The Reliability of Eyewitness Testimony*. Thesis Liberty University 440.

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

Sargent v. Mass. Accident Co., 29 N.E.2d 825 (Mass. 1940).

Simon, R. J., & Mahan, L. (1971). Quantifying Burdens of Proof: A View from the Bench, the Jury, and the Classroom. *Law & Society Review*, *5*(3), 319–330

Smith, M. (2018). When Does Evidence Suffice for Conviction? *Mind* 127 (508):1193-1218.

Smith v. Rapid Transit, Inc., 58 N.E.2d 754 (Mass. 1945).

Spottswood, M. (2016). Unraveling the Conjunction Paradox. *Law, Probability, and Risk*, 15(4): 259-296.

Thomson, J.J. (1986). Liability and Individualized Evidence. *Law and Contemporary Problems* 49(3): 199-219.

Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.

* Appendix & Section 5's integration into the dissertation project

Cases like Framing Twins suggest the need for an internalist condition on legal standards of proof. It could be that, in contrast to what is suggested in the concluding remarks of section 5, the satisfaction of legal standards of proof may require both an internalist and an externalist condition, just as is the case with some other epistemic standards (such as knowledge). Given this view, an internalist condition would merely state that meeting a standard of proof is partly a function of (at least) the three elements that Haack considers constitutive of the quality of the evidence for *p*. If this is the case, the modal view could be supplemented with an internalist condition, but Framing Twins would still constitute a counterexample to it insofar as it is still an easy possibility that Thelma is wrongfully held liable.

In contrast to the modal view, the knowledge-first probabilistic view would satisfy both internalist and externalist conditions (understood as above). However, it would still face cases like Misleading Paint in which $P(q|e) \leq .5$, unless KPE is amended so that the probability of the court's knowledge is calculated not given the evidence admitted in court but given the entire body of information of the court, *b*. If amended in this way, then plausibly $P(q|b) > .5$. This is because Gettier scenarios are themselves highly unlikely, so one does not normally consider the possibility of being in one given one's background knowledge. In this way, our background knowledge renders these scenarios unlikely, even though the admitted evidence itself may be silent about them. Such an amendment of KPE is reasonable and avoids the problem raised in Section 5.

By analyzing the modal view and knowledge-first probabilistic view, however, it is not clear whether an externalist condition is even necessary in an account of a standard of proof. This is so especially since even if an internalist component is added to MC, there will be counterexamples such as Framing Twins so long as the externalist condition is maintained. The causal condition, for example, is also externalist and is also too strong due to its externalist component (see section 5.2). Thus, inquiring over whether an externalist condition is necessary to meet a standard of proof is a larger, interesting project to pursue.

6. Concluding Remarks

This dissertation focused on the status of lottery propositions in relation to the epistemic standards of knowledge and justification, with the primary objective of seeking explanations of why we fail to know lottery propositions. It also discussed the status of litigated claims supported by bare statistical evidence –the legal correlates of lottery propositions– in relation to standards of proof.

Any explanation for why we fail to know lottery propositions will either i) reject that lottery beliefs are justified or ii) defend that lottery beliefs are true merely by epistemic luck. As argued in this dissertation, an explanation that carries out (ii) by providing an account of (veritic) epistemic luck as unsafe true belief fails. Consequently, we either renounce the thesis that lottery beliefs are epistemically luckily true or renounce the thesis that epistemic luck is unsafe true belief (and continue seeking accounts of epistemic luck that may explain why lottery beliefs are epistemically luckily true).

An explanation for why we do not know lottery propositions due to lottery beliefs being unjustified will identify either a probabilistic or a non-probabilistic condition on justification that such beliefs do not satisfy. As argued in this dissertation, and contrary to what Douven and Williamson (2006) argue, justification *can* have a probabilistic condition that lottery beliefs fail to satisfy. Probabilistic accounts of justification are thus restored as a theoretical possibility to solve the lottery paradox by rejecting that lottery propositions are justified.

If justification has a non-probabilistic condition that lottery beliefs do not satisfy, the best among four condition candidates is Martin Smith's (2016) normic support condition on justification. As argued in this dissertation, all conditions except the normic support condition face important problems that renders them unsuitable to explain why we do not know lottery propositions.

Alternatively, if lottery propositions are justified, there is nonetheless a solution to the lottery paradox compatible with their possessing such epistemic status. The solution consists in rejecting the principle of aggregativity of justification. As argued in this dissertation, Aggregativity is incorrect, and preserving this principle obstructs solving the most basic form of the lottery paradox. Of course, nothing prevents us from rejecting both aggregativity and the thesis that lottery propositions are justified.

Now, with respect to the legal correlates of lottery propositions –i.e., litigated claims supported by bare statistical evidence–, I divided explanations of why such claims do not meet a standard of proof into two types: externalist and internalist. Externalism about standards of proof establishes a condition on a standard mostly independent of the quality of the admitted evidence for the litigated claim, and explains how claims supported by bare statistical evidence do not satisfy such condition. In this dissertation, two externalist accounts of standards of proof were evaluated: Pritchard's modal view (2018, 2022) and Blome-Tillmann's knowledge-first probabilistic view (2017). The results were that, while they provide an attractive explanation of the insufficiency of bare statistical data to satisfy a standard of proof, they both independently face serious problems. Furthermore, such problems are ultimately explained by the fact that they are both externalist. This finding justifies the need to evaluate externalist accounts *qua* externalist to determine if an externalist account is a viable solution to the "proof paradox".

7. References

Blome-Tillmann, M. (2017). 'More Likely Than Not' - Knowledge First and the Role of Statistical Evidence in Courts of Law. In Adam Carter, Emma Gordon & Benjamin Jarvis (eds.), *Knowledge First: Approaches in Epistemology and Mind*. Oxford, UK: Oxford University Press. pp. 278-292.

Blome-Tillmann, M. (2020). Non-Reductive Safety. *Belgrade Philosophical Annual* 33 (33):25-38.

Blome-Tillmann, M. (2020). Statistical evidence, normalcy, and the gatecrasher paradox. *Mind*, *129*(514), 563–578.

Bondy, P. (2013). How to Understand and Solve the Lottery Paradox. *Logos & Episteme*, *4*(3), 283–292.

BonJour, L. (1985). *The structure of empirical knowledge*. Harvard University Press.

Buchak, L. (2014). Belief, Credence, and Norms. *Philosophical Studies*, 169(2): 285– 311.

Carnap, R., & Jeffrey, R. C. (1980). *Studies in inductive logic and probability* (Vol. 2). Univ of California Press.

Cohen, L.J. (1977). *The probable and the provable*. Oxford: Clarendon Press.

Di Bello, M. (2013). *Statistics and Probability in Criminal Trials: The Good, the Bad and the Ugly*. Doctoral Dissertation, Stanford University.

Douven, I. (2002). A new solution to the paradoxes of rational acceptability. British Journal for the Philosophy of Science 53 (3):391-410.

Douven, I. (2010). The pragmatics of belief. *Journal of Pragmatics*, 42(1), 35–47.

Douven, I. (2012). The lottery paradox and the pragmatics of belief. *Dialectica*, *66*(3), 351–373.

Douven, I., & Williamson, T. (2006). Generalizing the lottery paradox. *The British Journal for the Philosophy of Science*, *57*(4), 755–779.

Ebert, P., Durbach, I., & Smith, M. (2020). Varieties of risk. *Philosophy and Phenomenological Research*, 101, 432–455.

Engel Jr., M. (2024). Epistemic Luck. In The Internet Encyclopedia of Philosophy. https://iep.utm.edu/epi-luck/

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.

Feldman, R. & Conee, E. (2001). Internalism Defended. *American Philosophical Quarterly* 38 (1):1 - 18.

Fields, K., L. (2013). Toward a Bayesian Analysis of Recanted Eyewitness Identification Testimony. *New York University Law Review* 88(1769): 1769-1801

Foley, R. (1979). Justified inconsistent beliefs. *American Philosophical Quarterly*, *16*(4), 247–257.

Foley, R. (1987). *The theory of epistemic rationality*. Cambridge, MA: Harvard University Press.

Fratantonio, G. (2021). Evidence risk, and proof paradoxes: Pessimism about the epistemic project. *International Journal of Evidence and Proof*, 25, 307–325.

Gardiner, G. (2018). Legal Burdens of Proof and Statistical Evidence. In David Coady & James Chase (eds.), *The Routledge Handbook of Applied Epistemology*. New York: Routledge.

Gardiner, G. (2020). Relevance and risk: How the relevant alternatives framework models the epistemology of risk. *Synthese* 199 (1-2):481-511.

Gerken, M. 2011. Warrant and Action. *Synthese* 178: 529–47.

Haack, S. (2014). *Evidence Matters: Science, Proof, and Truth in Law*. Cambridge University Press.

Harman, G. (1968) Knowledge, Inference, and Explanation. *American Philosophical Quarterly*, *5*, 164-173.

Hawthorne, J. P. (2004). *Knowledge and lotteries*. Oxford University Press.

Hawthorne, J. and Stanley, J. (2008). Knowledge and Action. *Journal of Philosophy* 105: 571–90.

Hiller, A. & Neta, R. (2007). Safety and epistemic luck. *Synthese* 158 (3):303 - 313.

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

Ho, H.L. (2021). The Legal Concept of Evidence. *Stanford Encyclopedia of Philosophy* (Winter 2021 Edition), Edward N. Zalta (ed.), URL: < https://plato.stanford.edu/archives/win2021/entries/evidence-legal/>.

Ichikawa, J. J. (2014). Justification is potential knowledge. *Canadian Journal of Philosophy*, *44*(2), 184–206.

Kelly, K.T. & Lin, H. (2012). A Geo-logical Solution to the Lottery Paradox, with Applications to Nonmonotonic Logic. *Synthese* 186 (2):531-575.

Kelp, C. (2014). Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51, 227–232.

Kelp, C. (2016). Justified Belief: Knowledge First-Style. *Philosophy and Phenomenological Research*, *93*(1), 79–100.

Kelp, C. (2017a). Lotteries and justification. *Synthese*, 194(4), 1233-1244.

Kelp, C. (2017b). Knowledge first virtue epistemology. *Knowledge first: Approaches in epistemology and mind*, 223-45.

Kelp, C. (2019). How to be a reliabilist. *Philosophy and Phenomenological Research*, 98(2), 346-374.

Koehler, J. J., & Shaviro, D.N. (1990). Veridical Verdicts: Increasing Verdict Accuracy Through the Use of Overtly Probabilistic Evidence and Methods. *Cornell Law Review*, 75(2): 247–279.

Kvanvig, J.L. (2003). The Value of Knowledge and the Pursuit of Understanding. Cambridge University Press.Kyburg, H. E. (1961). *Probability and the logic of rational belief*. Middletown, CT, USA: Wesleyan University Press.

Laudan, L. (2006*). Truth, Error, and Criminal Law: An Essay in Legal Epistemology*. Cambridge: Cambridge University Press.

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

McCauliff, C.M.A. (1982). Burdens of Proof: Degrees of Belief, Quanta of Evidence, or Constitutional Guarantees?, 35 *Vanderbilt Law Review* 1293.

McGlynn, A. (2012). Justification as 'Would-Be' Knowledge. Episteme 9 (4):361-376.

McGlynn, A. (2014). *Knowledge First?* New York: Palgrave Macmillan.

Moser, P. K., & Tlumak, J. (1985). Two paradoxes of rational acceptance. *Erkenntnis*, *23*(2), 127–141.

Moss, S. (2018). *Probabilistic Knowledge*. Oxford: Oxford University Press.

Nance, D. (2021). Weight of Evidence. In *Philosophical Foundations of Evidence Law*. Oxford University Press.

Nelkin, D. K. (2000). The lottery paradox, knowledge, and rationality. *Philosophical Review*, *109*(3), 373–409.

Nelkin, D. (2021). Rational Belief and Statistical Evidence: Blame, Bias and the Law. In *Lotteries, Knowledge, and Rational Belief*. Cambridge University Press.

Nesson, C. (1979) Reasonable doubt and permissive inference—the value of complexity. *Harvard Law Review* 92: 1187–1199.

Ochoa, A. (2021). The Aggregativity of Justification and De-Generalizing the Lottery Paradox. [Unpublished Manuscript]

Pardo, M.S. (2019). The Paradoxes of Legal Proof: A Critical Guide. *Boston University Law Review*, 99.

Pollock, J. L. (1995). *Cognitive carpentry: A blueprint for how to build a person*. Mit Press.

Pollock, J. L. (2010). Defeasible reasoning and degrees of justification. *Argument and Computation*, *1*(1), 7–22.

Pritchard, D. (2005). *Epistemic luck*. Clarendon Press.

Pritchard, D. (2007). Anti-luck epistemology. *Synthese*, *158*(3), 277–297.

Pritchard, D. (2008a). Sensitivity, safety, and anti-luck epistemology. *The Oxford Handbook of Skepticism*, 437–455.

Pritchard D. (2008b). Knowledge, Luck and Lotteries. In V. Hendricks (ed.), *New Waves in Epistemology*. Basingstoke: Palgrave-Macmillan.

Pritchard D. (2012). Anti-Luck Virtue Epistemology. *Journal of Philosophy* 109(3), 247–279.

Pritchard D. (2015). Anti-luck Epistemology and the Gettier Problem. *Philosophical Studies* 172(1), 93–111.

Pritchard, D. (2015). Risk. *Metaphilosophy* 46 (3):436-461.

Pritchard, D. (2018). Legal risk, legal evidence and the arithmetic of criminal justice. *Jurisprudence* 9 (1):108-119.

Pritchard, D. (2022). In defence of the modal account of legal risk. *Synthese* 200 (4):1-16.

Rabinowitz, D. (2021). The Safety Condition for Knowledge. In *Internet Encyclopedia of Philosophy*. https://iep.utm.edu/safety-c/#SH2b

Redmayne, M. (2008). Exploring the proof paradoxes. *Legal Theory* 14 (4):281-309.

Roberts, N. (2014). *The Reliability of Eyewitness Testimony*. Thesis Liberty University 440.

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

Ryan, S. (1996). The epistemic virtues of consistency. *Synthese*, *109*(2), 121–141.

Salmon, W. C. (1963). *Logic* (Vol. 29, Issue 2). Englewood Cliffs, N.J., Prentice-Hall.

Sargent v. Mass. Accident Co., 29 N.E.2d 825 (Mass. 1940).

Simon, R. J., & Mahan, L. (1971). Quantifying Burdens of Proof: A View from the Bench, the Jury, and the Classroom. *Law & Society Review*, *5*(3), 319–330

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, 2016, 213 pages*.

Smith, M. (2018). When Does Evidence Suffice for Conviction? *Mind* 127 (508):1193-1218.

Smith v. Rapid Transit, Inc., 58 N.E.2d 754 (Mass. 1945).

Spottswood, M. (2016). Unraveling the Conjunction Paradox. *Law, Probability, and Risk*, 15(4): 259-296.

Sosa, E. (1999). How to Defeat Opposition to Moore. *Philosophical Perspectives* 13, 137-149.

Sosa, E. (2010). How Competence Matters in Epistemology. *Philosophical Perspectives*, 24, 465-475.

Silva, P. & Oliveira, L.R.G. (2024). Propositional Justification and Doxastic Justification. In Maria Lasonen-Aarnio & Clayton Littlejohn (eds.), *The Routledge Handbook of the Philosophy of Evidence*. New York, NY: Routledge.

Sylvan, K. (2014). *On Divorcing the Rational and the Justified in Epistemology*. https://philarchive.org/archive/SYLODT

Talbott, W. (2016). Bayesian Epistemology. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/win2016/entries/epistemology-bayesian/.

Thomson, J.J. (1986). Liability and Individualized Evidence. *Law and Contemporary Problems* 49(3): 199-219.

Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.

Yang, D. (2019). What's Wrong with Modal Conceptions of Luck and Risk. *Erkenntnis* 86(4), 773–787.