




ARTICLE

# Probabilistic Truth, Relativism, and Objective Chance

Svenja Schimmelfennig 

University of Cologne, Cologne, Germany  
Email: [schimmelfennig.svenja@gmail.com](mailto:schimmelfennig.svenja@gmail.com)

(Received 20 September 2022; revised 21 November 2022; accepted 20 December 2022;  
first published online 20 February 2023)

## Abstract

In *Probabilistic Knowledge* Sarah Moss proposes that our credences and subjective probability judgments (SPJs) can constitute knowledge. Mossean probabilistic knowledge is grounded in probabilistic beliefs that are justified, true, and unGettiered. In this paper I aim to address and solve two challenges that arise in the vicinity of the factivity condition for probabilistic knowledge: *the factivity challenge* and *the challenge from probabilistic arguments from ignorance* (probabilistic AIs). I argue that while Moss's deflationary solution to the factivity challenge formally works, it leaves us ill-equipped to handle probabilistic AIs. An account of probabilistic knowledge that cannot overcome probabilistic AIs makes knowledge of thoroughly probabilistic contents a rare and unstable phenomenon, at best, or, at worst, impossible. I hold that establishing a metaphysically enriched account of probabilistic truth is therefore mandatory. I go on to develop a truth-conditional approach on probabilistic factivity that is relativistic in its nature and centers on objective chances. I show that while the approach is still compatible with Moss's overall semantics for probabilistic knowledge, it provides us with a simple but forceful answer to probabilistic AIs.

**Keywords:** Probabilistic truth; probabilistic knowledge; truth-relativism; objective chance; probabilistic Gettier cases

## Introduction

In *Probabilistic Knowledge* Sarah Moss (2018b) proposes that our credences and subjective probability judgements (SPJs) can constitute knowledge. The idea that there is probabilistic knowledge is compelling in many ways. If there is such a thing as probabilistic knowledge, it would allow us to build connections between formal and traditional epistemology, discuss probabilistic knowledge norms for credences and action, as well as uncover novel connections between statistical evidence, credence, and knowledge. Mossean probabilistic knowledge is grounded in probabilistic beliefs that are justified, true, and unGettiered. In this paper I aim to address and solve two challenges that arise in the vicinity of the factivity condition for probabilistic knowledge. The first challenge consists in overcoming the worry that credences and SPJs might not be truth-apt. The proponent of probabilistic knowledge has to provide a convincing account that

© The Author(s), 2023. Published by Cambridge University Press

explains how and by virtue of what probabilistic beliefs can be true. Following Rich (2020), call this the *factivity challenge*. The second challenge arises from probabilistic arguments from ignorance (probabilistic AIs). Probabilistic AIs challenge the justificatory status of a lot of our everyday probabilistic beliefs by calling our attention to non-probabilistic alternatives that we cannot rule out on statistical evidence alone. Though probabilistic AIs, therefore, primarily put pressure on probabilistic knowledge via the justification condition, I hold that both challenges are connected. Different ways of solving the factivity challenge equip us more or less well for dealing with probabilistic AIs. I argue that while Moss's deflationary solution to the factivity challenge formally works, it leaves her ill-equipped to handle probabilistic AIs. An account of probabilistic knowledge that cannot overcome probabilistic AIs makes knowledge of thoroughly probabilistic contents a rare and unstable phenomenon, at best, or, at worst, impossible. I hold that establishing a metaphysically enriched account of probabilistic truth is therefore mandatory. After introducing the factivity challenge and the challenge from probabilistic AIs in §1 and §2, I demonstrate in §3 that adopting a version of truth-relativism that centers on epistemic probabilities is adequate for overcoming both challenges. However, as I show in §4, a version of truth-relativism that centers on epistemic probabilities cannot account for the correct verdicts concerning probabilistic Gettier cases. I therefore use §5 to develop another version of truth-relativism that centers on objective chances rather than epistemic probabilities. I argue that this relativist approach on probabilistic truth allows us to overcome the factivity challenge and probabilistic AIs in the same way as the more familiar epistemic probability relativism, while also yielding the correct verdicts in probabilistic Gettier cases.

## 1. The Factivity Challenge

According to Moss, credences and SPJs are full beliefs in complex probabilistic contents (as opposed to complex attitudes with simple propositional contents).<sup>1</sup> Probabilistic contents, in turn, are conceived as sets of probability spaces. Formally speaking, a probability space is an ordered triple,  $s = \langle \Omega_s, \mathcal{F}_s, \mu_s \rangle$ , containing a domain of possibilities,  $\Omega_s$ , constituted by a non-empty set of possible worlds, an event space,  $\mathcal{F}_s$ , which is a set of subsets of  $\Omega_s$  (and thereby a set of propositions), and a probability function,  $\mu_s$ , that assigns each proposition in  $\mathcal{F}_s$  a value between 0 and 1.<sup>2</sup> To provide a quick example, the content that Jones probably smokes is representable as a set of probability spaces over a domain of possibilities containing two worlds: one where Jones smokes and one where she does not. The set contains each probability space that assigns a value above 0.5 (including 1) to the proposition that Jones smokes. All other possible probability spaces are excluded from the set.<sup>3</sup>

Conceiving of probabilistic beliefs as full beliefs in probabilistic contents leads Moss (2018b: 86) to conclude that for a probabilistic belief to constitute knowledge means “just the same as whatever it traditionally means to say that beliefs constitute knowledge.” While it seems unproblematic to hold that credences and SPJs can be rational and/or justified, the truth-aptness of probabilistic beliefs is often questioned.<sup>4</sup>

<sup>1</sup>See Moss (2018b: 1–19) for arguments in favor of the complex content over the complex attitude view.

<sup>2</sup>See Moss (2018b: 231), Demey *et al.* (2019), and Rich (2020: 1709).

<sup>3</sup>Throughout this paper I am using Kratzer's (1991) semantic interpretation of ‘probably’ which is also followed by Moss (2018b: 85). See Rich (2020: 1709) for another example.

<sup>4</sup>See Moss (2018b: 120).

Knowledge is factive. If we have knowledge of any content  $\psi$ , it follows that  $\psi$  is true. A proponent of Mossean probabilistic knowledge has to show that probabilistic contents are truth-apt and, therefore, provide a theory of probabilistic truth. The factivity challenge for probabilistic knowledge consists in the obstacle of providing a plausible account of probabilistic truth.

In *Probabilistic Knowledge* Moss introduces a deflationary solution to the factivity challenge according to which a probabilistic content  $\langle \psi \rangle$  is true iff  $\psi$ .<sup>5</sup> While deflationism formally solves the factivity challenge, it is frequently criticized for leaving probabilistic truth, and therefore probabilistic knowledge, a concept way too opaque to comfortably operate with.<sup>6</sup> I hold that this opacity is not the most pressing issue for the deflationary approach. Above all, Moss's deflationism leaves her ill-equipped to properly handle probabilistic AIs and thereby unable to defend the concept of probabilistic knowledge against skepticism.

## 2. The Challenge from Probabilistic Arguments from Ignorance

Without further ado consider the following standard lottery case:

### Lottery

Thilo has purchased a ticket in a 1000 ticket lottery with exactly one winner. He knows these facts about the lottery. Based on the statistical evidence Thilo forms the thoroughly probabilistic belief that his ticket has probably lost. Does he know this?

Intuitively, we would like to judge that Thilo knows that his ticket has probably lost.<sup>7</sup> Let us assume, for the moment, that it is indeed the case that his ticket has probably lost. This suffices to satisfy Moss's deflationary definition of probabilistic truth. However, the claim that Thilo knows that his ticket probably lost is challenged by the following probabilistic argument from ignorance:

(AI1) If Thilo knows that his ticket has probably lost, he is in a position to know that it is not the case that his ticket has certainly won.

(AI2) Thilo is not in a position to know that it is not the case that his ticket has certainly won.

(AI3) Therefore, Thilo does not know that his ticket has probably lost.

Let us consider each premise individually. The epistemic principle of K-closure underlying AI1 is *prima facie* plausible.<sup>8</sup> It seems that if I know that  $p$ , and that  $p$  entails  $q$ , I

<sup>5</sup>See Moss (2018b: 120–1). See also Blackburn (1984) and Gibbard (2003) for defenses of the deflationary theory of truth in the context of ethical expressivism.

<sup>6</sup>See for instance Easwaran (2018), Smartt (2019), and Rich (2020).

<sup>7</sup>It is important to notice that there are multiple readings of sentences containing epistemic vocabulary. The content of 'Thilo probably lost' can be either conceived as a propositional content *about* probabilities, or as a thoroughly probabilistic content. In the context of this paper, sentences containing epistemic vocabulary are always intended to be understood in the second way (if not indicated otherwise). See also Moss (2018b: 86).

<sup>8</sup>DeRose (1995: 26–7), for instance, argues that in denying K-closure we are forced to embrace the correctness of "abominable conjunctions" of the following kind: "I cannot know that  $q$ , but still know that  $p$  and that  $p$  implies that  $q$ ." Even if we deny K-Closure, it is not clear that such a move is helpful in

am, at least, in a position to know  $q$ . The content of ‘Thilo has probably lost’ ( $\psi$ ) and the content of ‘he has certainly won’ ( $\varphi$ ) are inconsistent. If Thilo has certainly won, it is false that he has probably lost.<sup>9</sup> Hence, if Thilo knows that he has probably lost (and recognizes the above inconsistency) he should be able to know that it is false that he has certainly won.

This takes us directly to AI2 according to which Thilo cannot not know that it is not the case that his ticket has certainly won. Following Moss, I call the content that Thilo has *certainly* won a *nominally probabilistic content*. The content of ‘Thilo has certainly won’ is a set of probability spaces where each space has only worlds in its domain in which Thilo actually wins. The content of ‘Thilo probably lost’ is different. It is *thoroughly probabilistic*, for it is represented by a set of probability spaces some of which do have worlds in their domain in which Thilo actually wins.<sup>10</sup> On Moss’s semantics nominally probabilistic contents are type-shifted versions of simple propositional contents. This semantic approach is motivated by the fact that we often infer full beliefs from credences and vice versa. While Moss’s account captures the logical connections between simple and probabilistic contents, disjunctive theories fail to account for the validity of arguments featuring both kinds of contents.<sup>11</sup> This motivates a one-to-one correspondence between propositions and nominally probabilistic contents. Believing that  $p$  is then, strictly speaking, nothing more than having the nominally probabilistic belief that *certainly*  $p$ .<sup>12</sup>

With this in mind we should be able to see why AI2 is hard to deny. It is standardly assumed that in lottery cases we are not in a position to know that it is not the case that we have (certainly) won, even though the chances of (certainly) winning are extremely low. It is widely accepted that having merely statistical evidence for  $p$  does not suffice for us to come to know  $p$ .<sup>13</sup> This explains why people regularly participate in lotteries with low chances of winning. After all, you never know.

AI1 is plausible on its own merits and AI2 is hard to deny on Moss’s semantics. Moss indeed accepts both premises. If we do so, AI3 follows directly. Since probabilistic arguments from ignorance can be construed for a wide range of our credences and SPJs,

---

overcoming probabilistic AIs. Consider, for instance, Hawthorne (2005) who argues that K-Closure does not apply when “manifestly heavyweight” propositions are involved. Heavyweight propositions are propositions for which it is hard to see how they could ever be known (e.g. skeptical hypotheses). However, it is far from obvious that the proposition that Thilo has (certainly) won is such a heavyweight proposition. After all, it appears to be rather lightweight. Additionally, Moss (2018a, 2021) relies on closure in her arguments for moral encroachment on credences.

<sup>9</sup>Formally this can be derived from the fact that  $\psi$  and  $\varphi$  are disjoint sets of probability spaces. See Moss (2018b: 53–4).

<sup>10</sup>In Moss’s (2018b: 14) words: nominally probabilistic contents merely “represent ... some distinction between possible worlds, namely those that are in their domain and those that are not.”

<sup>11</sup>See Moss (2018b: 53–7) for examples.

<sup>12</sup>One might object that we often seem to believe  $p$  without believing *certainly*  $p$ . To alleviate the worry Moss (2018b: 58–63) introduces a difference between strict and loose contents of simple beliefs. She argues that when we hold or express the simple belief that  $p$ , we believe or express the strict content that *certainly*  $p$  plus a different thoroughly probabilistic loose content that depends on the given pragmatic or contextual conditions. This loose content could, for instance, be that *almost certainly*  $p$ . See Moss (2019) for further reasons and advantages of interpreting simple sentences and simple belief ascriptions as loose speech.

<sup>13</sup>Turri (2021) provides empirical evidence for our reluctance to attribute knowledge in standard lottery cases. Pritchard (2008) derives the conclusion from the anti-luck platitude for knowledge, while Nelkin (2000) gives an argument along the lines of Kyburg’s (1961) lottery paradox. Williamson (2002: 246–52) argues that the infelicitousnesses of lottery assertions is best explained by the fact that lottery propositions cannot be known.

namely against all probabilistic beliefs that are held on the grounds of merely statistical evidence, they call into question the possibility of probabilistic knowledge in a wide variety of cases. This seems false. If probabilistic knowledge, as Moss presents it, is supposed to be a plausible and fruitful concept for epistemologists, her account should allow for probabilistic knowledge in cases where we expect it the most: lotteries, coin-tosses, and, more broadly, in cases where our evidence does not suffice to justify full beliefs, but arguably credences.

Moss indeed takes the challenge from probabilistic AIs seriously and proposes a solution that seemingly allows her to salvage probabilistic knowledge without having to deny any of the premises A1–A3. Adopting relevant alternatives contextualism (RAC) for (probabilistic) knowledge, she argues that probabilistic AIs are such that they create contexts in which challenging nominally probabilistic contents become relevant. Since the RAC theorist holds that a subject *S* knows *p* iff *S* truly believes *p* on the basis of evidence that eliminates all *relevant alternatives* to *p*<sup>14</sup>, in the context of the probabilistic AI Thilo will indeed lack probabilistic knowledge. However, this does not imply that Thilo could *never* know that his ticket probably lost. Moss holds that at most *ordinary* contexts nominally probabilistic alternatives are not relevant and therefore need not be ruled out in order to arrive at probabilistic knowledge.<sup>15</sup>

However, issues are lurking. Moss seems to be convinced that by raising a nominally probabilistic alternative against a thoroughly probabilistic belief we create contexts of a similar kind as by raising brain-in-a-vat-like (BIV-like) alternatives to ordinary full beliefs. But, intuitively, nominally probabilistic contents – like the one that Thilo’s ticket has certainly won – are not as aberrant as the skeptical alternatives raised in standard AIs.

Proponents of RAC provide different principles for determining whether an alternative counts as relevant at a given context. Moss explicitly accepts the following criterion by Lewis (1996: 559):

**Rule of Attention (ATT)**

If you pay attention to some alternative at a given context, then the alternative will be relevant for you at this context.

ATT states that we cannot properly ignore an alternative if we actively consider the alternative. Only because a subject is ignorant of a possibility it does not follow that she is in a position to *properly* ignore it. In addition to ATT we are, therefore, in need of further criteria to determine when alternatives are relevant. Consider two further rules given by Lewis (1996: 554, 556):

**Rule of Actuality (ACT)**

If some content is true, then it is relevant as an alternative.

**Rule of Resemblance (RES)**

Suppose one possibility saliently resembles another. Then if one of them may not be properly ignored, neither may the other.

ACT and RES in combination yield that if, at context *C*, a true (and thereby relevant) alternative saliently resembles another possibility, this possibility will be relevant no

<sup>14</sup>See Goldman (1988: 197).

<sup>15</sup>See Moss (2018b: 133–7).

matter whether it is ignored or false at C. Of course these rules only give us a vague guidance as to when alternatives are relevant. Nonetheless, together with the above observation that the nominally probabilistic alternatives in probabilistic AIs are way less aberrant than the alternatives raised by standard skeptical AIs they already suffice to uncover some pressing issues for Moss's RAC approach.

Let us focus on ATT first. It seems that while a proponent of RAC for simple knowledge claims can properly claim that in most ordinary contexts (outside the philosophy classroom) I do not pay attention to the possibility of being a BIV, the same does not obviously hold for nominally probabilistic alternatives that are raised against thoroughly probabilistic beliefs in probabilistic AIs. Consider again Lottery. Thilo knows how lotteries work. He is aware that winning is a possibility. Hence, he also is likely to be aware that it is possible that his ticket has (certainly) won instead of probably lost. This indicates that nominally probabilistic alternatives of the above kind are ordinarily and frequently relevant and, therefore, very different from other skeptical possibilities. If it is true that we ordinarily do pay attention to nominally probabilistic alternatives, we will frequently fail to have knowledge of the most ordinary probabilistic contents.<sup>16</sup>

Despite the above worry, let us assume that there are contexts where Thilo is not aware of the possibility that his ticket has (certainly) won. Let us further assume that Thilo believes that his ticket has probably lost. According to Moss, Thilo will be able to know this. He is in a position to rule out other relevant thoroughly probabilistic alternatives (e.g. that it is exactly .4 likely that his ticket has lost) on his evidence and he does not need to rule out the possibility of (certainly) winning because it is not relevant at the context of his belief. But now imagine that Thilo transforms his initial probabilistic belief into a slightly stronger one, namely the belief that his ticket has certainly lost. We can imagine that nothing further changes. Thilo continues to be ignorant of the possibility of his ticket being the winner. It seems that we would have to conclude that Thilo knows that he has certainly lost. But this is obviously false. In fact, ACT in combination with RES explain this verdict. Suppose that ticket #37 wins. According to ACT the possibility of ticket #37 winning is thus relevant. Even if Thilo is not the holder of ticket #37, RES tells us that, since the possibility of Thilo's ticket winning saliently resembles the relevant alternative of ticket #37 winning, Thilo's ticket winning is a relevant alternative as well. Thilo cannot rule out that his ticket has (certainly) won on his evidence; therefore he cannot know that his ticket has (certainly) lost.<sup>17</sup>

If the nominally probabilistic alternative *is* relevant for Thilo according to ACT and RES at the ordinary context of Thilo's nominally probabilistic belief, why should it not be relevant at the ordinary context of Thilo's previously held thoroughly probabilistic belief? This is what I call the asymmetry challenge. The proponent of RAC as a solution to probabilistic AIs has to give an explanation as to why switching between thoroughly and corresponding nominally probabilistic beliefs has an impact on context such that in the latter case nominally probabilistic alternatives are naturally relevant while in the former case they are not. As long as such an asymmetry cannot be made plausible, Moss faces a dilemma. Either she has to conclude that we are sometimes able to know on the

<sup>16</sup>It might be argued that this is just a straightforward consequence of RAC in general. Proponents of RAC are happy to admit that the moment I consider that I am a BIV, I will lose a lot of ordinary knowledge. Once I shift my attention away from the skeptical hypothesis I regain my knowledge. But notice that while we often find ourselves ignoring BIV-like hypotheses, it is rare that we ignore such simple possibilities as winning the lottery.

<sup>17</sup>See Lewis (1996: 557).

basis of merely statistical evidence that our lottery ticket has (certainly) lost, or we have to admit that whenever it is impossible for us to know a proposition  $p$  on the basis of merely statistical evidence, it will also be impossible for us to have probabilistic knowledge of  $p$  (e.g. knowledge that *probably*  $p$ ).

To be sure, I am not denying that changing the content of one's belief *can* invoke shifts in context that could make certain previously irrelevant alternatives relevant. Assume, for instance, that Thilo fully believes that he has hands at an ordinary context where he does not consider any skeptical hypothesis. Now imagine that Thilo moves from believing that he has hands to believing that he is not a BIV. Obviously, this change of belief also invokes a broader contextual shift. By believing that he is not a BIV Thilo considers the possibility of being a BIV. He does not hold this belief at an ordinary but rather at a skeptical context. However, this mechanism does not obviously apply for the change from believing a thoroughly to believing a corresponding nominally probabilistic content. Remember that in order for RAC to deliver a proper defense of probabilistic knowledge against skepticism, we have to assume that we can (at least sometimes) properly ignore challenging nominally probabilistic contents. It is not clear at all why Thilo should be able to properly ignore the possibility of his ticket winning while he believes that it has probably lost, but cease to be able to do so as soon as he starts believing that his ticket has (certainly) lost.

Though Moss does not directly discuss the asymmetry challenge she defends the following rule

#### **General Asymmetry Rule (GAR)**

Thoroughly probabilistic beliefs are more easily challenged by other thoroughly probabilistic contents, and less easily challenged by nominally probabilistic contents.<sup>18</sup>

We can charitably assume that if GAR is true, it will also turn out to be true that nominally probabilistic contents are more challenging for other nominally probabilistic contents than for thoroughly probabilistic ones. Anyways, I hold that Moss's arguments for GAR fail. Moss sketches two possible reasons to assume that GAR is correct. Let us consider each of them in turn.

The first reason Moss cites in favor of GAR rests upon the idea that certain nominally probabilistic contents are 'routinely present' when we hold thoroughly probabilistic beliefs. When Thilo believes that his ticket probably lost, he necessarily believes that his ticket *might* be the winning ticket. According to Moss, this makes the alternative that his ticket has (certainly) won routinely present for him and thereby less challenging for his probabilistic belief. Other thoroughly probabilistic alternatives are not routinely present for Thilo and appear more challenging for this reason. Extending this reasoning, we could argue that when Thilo believes that his ticket (certainly) lost he does not believe that his ticket might have won. His belief is thereby more easily challenged by the alternative of (certainly) winning.

I hold, however, that this approach is flawed. First of all, note that Moss's introduction of *challenging* alternatives is infelicitous. It masks what kind of asymmetry GAR really has to account for. As I see it, some content  $\psi$  is challenging for some other content  $\varphi$  iff  $\psi$  is inconsistent with  $\varphi$  and relevant at the context of  $\varphi$ 's use. Furthermore, if  $\psi$  is inconsistent with  $\varphi$ ,  $\psi$ , by definition, constitutes an alternative to  $\varphi$ . In this light,

<sup>18</sup>See Moss (2018b: 138).

the possibilities introduced by probabilistic AIs are certainly *alternatives*. Whether they challenge the content in question does then wholly depend on whether they are ordinarily *relevant*. This is why the first argument for GAR fails. It is just hard to see how something like the *routine presence* of a nominally probabilistic alternative at a context could make the alternative *less relevant* at this same context. Moss endorses ATT. If a possibility is routinely present for us it should be safe to say that we are routinely aware of it. With ATT it follows that the same possibility is routinely relevant for us. In this light, Moss's observation of nominally probabilistic alternatives being routinely present relative to our thoroughly probabilistic beliefs should leave us even more concerned about the possibility of acquiring and maintaining probabilistic knowledge rather than placate the worries raised by probabilistic AIs.

Moss offers a second argument for GAR. Following Weatherson (2011) she argues that there is a difference between AIs that raise merely abstract possibilities and those that make us aware of more ordinary possibilities that we have not ruled out yet, such that the former arguments appear to be less effective than the latter. Moss holds that raising nominally probabilistic alternatives against thoroughly probabilistic beliefs is a merely abstract move. Nominally probabilistic contents are therefore less challenging relative to thoroughly probabilistic beliefs. But this reasoning presupposes the required asymmetry without explaining it. Why are nominally probabilistic alternatives merely abstract possibilities relative to thoroughly probabilistic contents? As we have seen, the possibility of Thilo's ticket (certainly) winning does not appear to be skeptical, far-off, or abstract in any way similar to a BIV-hypothesis. Furthermore, the possibility of winning surely does not constitute an abstract alternative relative to the belief that Thilo (certainly) lost. So why should it be abstract in the context of believing that Thilo's ticket probably lost? We are returned to the initial asymmetry challenge.

I conclude that GAR is implausible. There are no obvious relevant contextual differences invoked by transitioning from believing a thoroughly probabilistic to believing a corresponding nominally probabilistic content. If the possibility of winning is relevant at the context of Thilo's belief that his ticket has (certainly) lost, it should also be relevant at the context of his belief that his ticket probably lost. ATT, as well as ACT and RES do not suggest that we should find any asymmetry here. Nonetheless, I am convinced that we do not have to bite the bullet and conclude that we (almost) never are able to obtain probabilistic knowledge on merely statistical evidence. A plausible story can be told about how ordinary nominally probabilistic alternatives are always relevant and yet can be ruled out on purely statistical evidence relative to *thoroughly* probabilistic beliefs. Such a story demands an inflated account of probabilistic truth. The general idea is the following: by saying more about probabilistic truth we are revealing what our probabilistic beliefs aim at. An inflation of probabilistic truth will thus uncover important asymmetries and provide us with a better grip on the concept of probabilistic knowledge. In the following section I introduce such inflation and argue that it is compatible with Moss's overall semantic theory of probabilistic contents, beliefs, and knowledge claims. I then show how the account allows us to meet the challenge posed by probabilistic AIs *without* running into the asymmetry challenge.

### 3. Inflating Probabilistic Truth

An important step towards exacting the concept of probabilistic truth is provided by Rich's (2020) introduction of Kripke models for probabilistic knowledge and, consequently, probabilistic factivity. Put briefly, Rich's model  $M = \langle \Omega, R, V^P \rangle$  is a tuple



consisting of a set of possible worlds,  $\Omega$ , a valuation function,  $V^P$ , mapping atomic propositions onto sets of possible worlds, and an accessibility relation,  $R$ , specifying which probability spaces the agent does not rule out at each world  $\omega \in \Omega$ .<sup>19</sup>

Let us define  $S(P)$  as the set of all probability spaces,  $s$ , for which the domain of possibilities is a subset of, or identical to,  $\Omega$  as defined in  $M$ . Let  $\psi \subseteq S(P)$ . Probabilistic truth is then modeled as follows:

$$(M, \psi) \models T\varphi \text{ iff } (M, \psi) \models \varphi \text{ where } (M, \psi) \models \varphi \text{ iff } \psi \subseteq \varphi^{20}$$

According to probabilistic Kripke models, the factivity of a probabilistic content  $\varphi$  is dependent on a probabilistic reference point  $\psi$ .  $\varphi$  is true relative to reference point  $\psi$  iff  $\psi$  implies  $\varphi$ .  $\psi$  implies  $\varphi$  iff it is a subset of, or equivalent to,  $\varphi$ . Saying more about probabilistic truth, therefore, amounts to determining conditions for choosing reference points.

Additionally, the model allows us to identify a first asymmetry between factivity for propositional and probabilistic contents. According to standard Kripke models a propositional content  $p$  is true *at a world*  $\omega$  iff  $\omega \in V(p)$ .<sup>21</sup> Since for Moss propositions are directly translatable into nominally probabilistic contents, it is natural to assume that nominally probabilistic contents are true relative to the probabilistic analog of a possible world. This analog is what Moss calls a *divine probability space*.<sup>22</sup> Divine probability spaces are probability spaces with only one world in their domain. The probabilistic equivalent of the actual world would thus be a probability space containing only the actual world. In this way there is still a pretty clear sense in which nominally probabilistic contents are true *at worlds*. Remember that the simple content that  $p$  can be displayed as a set of probability spaces,  $\varphi$ , according to which  $p$  is certain (i.e. true at every world of the domain of  $\varphi$ ). Whether  $p$  is true at  $\omega$  then depends on whether the probability space that has only  $\omega$  in its domain is a subset of  $\varphi$ . This, in turn, boils down to the question whether  $\omega$  is a member of the domain of  $\varphi$  which amounts to the standard Kripke interpretation of propositional factivity.

Thoroughly probabilistic contents do not correspond to simple propositions. If we were to treat them like nominally probabilistic contents and standardly assess them relative to divine probability spaces, they could only ever be trivially true. For instance, Thilo's belief that his ticket has probably lost could only be true if the world in which Thilo holds this belief would be such that Thilo's ticket indeed lost. As a further consequence, moderate credences could never be true and knowledge. Moss argues that we can know probabilistic contents as specific as 'Jones is .6 likely to smoke.' I hold that, quite intuitively, Thilo does not only know that his ticket has probably lost, but is also in a position to know that it is .999 likely that his ticket has lost.<sup>23</sup>

The question is whether we can find conditions to pick out reference points which allow for thoroughly probabilistic contents to be non-trivially true. Rich suggests three possible kinds of reference points that could play the desired role: objective chances, evidential probabilities, or unique rational credences. Furthermore she observes that the

<sup>19</sup>See Rich (2020: 1712–13).

<sup>20</sup>See Rich (2020: 1715).

<sup>21</sup>Where  $V$  is a function mapping  $p$  onto a set of possible worlds.

<sup>22</sup>See Moss (2018b: 134).

<sup>23</sup>Remember that the here intended reading is not propositional. The point is not that Thilo knows that the evidential or objective probabilities of losing are .999. It is about his .999 credence constituting knowledge.

correct choice of reference point is likely to depend (in some way or other) on contextual features.<sup>24</sup> With this in mind we can further specify the questions an inflated account of probabilistic truth for thoroughly probabilistic contents is supposed to answer:

1. Which features at which context are relevant for picking out reference points?
2. Of which kind are the reference points suggested by these features?

### 3.1. *Evidential Probability Relativism*

Let us begin by considering the first question. The idea that truth-conditions for sentences embedding epistemic modals are context-dependent is not surprising at all. As MacFarlane (2011: 144) puts it: “Clearly, epistemic modals have something to do with knowledge.” Whether it is true that Thilo’s ticket has probably lost depends partly on what is known about the issue and what is known is variable over contexts. This already equips us with a partial answer to the first question: the relevant features that are decisive for picking out reference points are epistemic. More precisely, the relevant feature appears to be someone’s total knowledge and/or evidence. Whose knowledge exactly? I take this question to be tantamount to the question of *which* context is relevant for picking out reference points. Roughly, two answers can be given: We can either focus on the *context of use* or the *context of assessment* of a given utterance or belief to provide us with the relevant set of known propositions. Focusing on a subject’s total knowledge at the context of use will amount to defending some version of contextual truth, while making truth depend on what is known at the context of assessment will make us proponents of some sort of truth-relativism for probabilistic contents. I hold that we ought to reject contextualism in favor of relativism.

The first reason for doing so is that a contextual account is not going to yield the right verdicts concerning third-party assessments, retractions, and disputes involving probabilistic contents. While a contextualist will typically claim that it is Thilo’s total knowledge that is decisive for the truth of his belief, a relativist is going to argue that it is all that is known by the assessor(s) of Thilo’s belief – in this case the readers – that is crucial. Now imagine that Thilo claims that he knows that his ticket has probably lost, while we already know that it has won. Surely, we would like to judge that Thilo is wrong. He *is* the lucky winner, and the lucky winner has not probably lost. A contextualist, however, wants us to respond in a very different way. She holds that for all Thilo knows it is still true that he probably lost. If this is correct, we should be happy to admit that Thilo justifiably believes something true and has, therefore, knowledge. But we do not appear willing to do so.<sup>25</sup> The relativist, on the other hand, can account for a genuine disagreement between us and Thilo. This is because, according to the relativist, it is *our* knowledge that determines the correct reference point according to which the truth of the content of Thilo’s belief is to be assessed. Given all that *we* know, Thilo has won and therefore not probably lost.<sup>26</sup>

Secondly, contextualism about epistemic modals is highly incompatible with Moss’s conception of probabilistic contents, beliefs, and knowledge. On the contextualist’s view

<sup>24</sup>See Rich (2020: 1723).

<sup>25</sup>Sometimes we concede things like “For all Thilo knows, it is true that he probably lost.” However, when we grant something like that we are not talking about probabilistic truth anymore. We are now concerned with a proposition about (Thilo’s) epistemic probabilities.

<sup>26</sup>See MacFarlane (2011: 146–68) for more detailed arguments against contextualism and for relativism.

the contents of sentences embedding epistemic modals are not sets of probability spaces, but sets of possible worlds. A solipsistic contextualist would interpret the content that Thilo's ticket probably lost as meaning that for all that Thilo knows it is the case that his ticket probably lost. This is a simple propositional content. As a consequence Thilo's belief would have to be interpreted as a standard full belief. It is true that there are generally two readings of sentences embedding probabilistic vocabulary: a probabilistic *and* a propositional one. Moss offers compelling reasons to assume that we are able to clearly distinguish the probabilistic and propositional reading.<sup>27</sup> The contextualist, however, carelessly conflates the two readings. Relativism, on the other hand, is compatible with Moss's overall theory of probabilistic knowledge. Though Moss points out that her "probabilistic theory of epistemic vocabulary does not make use of any relation of relative truth," she also concedes that her theory is "relativist friendly."<sup>28</sup> Truth-relativism can be conceived as a middle-ground between standard truth-conditional contextualist theories and non-truth-conditional expressivist theories about epistemic vocabulary. It acknowledges the problems contextualism faces without denying – as the expressivist does – that there are (non-deflationary) truth-conditions for epistemic modals. Furthermore, relativists about epistemic modals usually agree that the contents of probability statements are not sets of possible worlds, but more complex contents. For instance, on Egan's (2007: 5) and Stephenson's (2007: §4) preferred versions of relativism contents of probabilistic utterances and beliefs are sets of *centered* worlds (i.e. <world, time, individual> triples), while MacFarlane (2011: 166–8) conceives of probabilistic contents as <world, information state> tuples. An information state, in turn, is construed as the set of probability spaces that are not ruled out by a subject's total knowledge. MacFarlane's choice to use world-information-state-pairs is owed to the fact that he needs to allow for joint statements containing probabilistic and propositional contents to have truth values.<sup>29</sup> On Moss's approach such maneuvering is not necessary, for propositional contents simply are the limiting cases of probabilistic contents. As I see it, nothing prevents us from adopting a version of relativism maximally compatible with Moss's semantics where probabilistic contents are merely sets of probability spaces and probabilistic truth is dependent on a reference point determined by the total knowledge of the assessor.

Moving on to the second question an account of inflated probabilistic truth has to answer: Of which kind are the reference points that are picked out by the total knowledge of a subject at the context of assessment? Given what has been said so far, a natural suggestion is the following: Regarding the question of whether  $p$ , an assessor's total knowledge uniquely determines a probability space that assigns her evidential probabilities to  $p$ .<sup>30</sup> It is this reference point relative to which the truth of a probabilistic content concerning whether  $p$  is to be assessed. Call this kind of reference point an *evidential probability reference point* and the resulting version of truth-relativism *evidential probability relativism* (EPR).

<sup>27</sup>See Moss's (2018b: 33–7) test battery for probabilistic contents.

<sup>28</sup>Moss (2018b: 124, 128).

<sup>29</sup>See MacFarlane (2019: 98–100).

<sup>30</sup>These evidential probabilities could, for instance, be interpreted as the credences an ideally rational agent possessing the same total evidence as the assessor would adopt. See Eder (2019) for defense of this interpretation of evidential probabilities against objections by Williamson (2002).

### 3.2. Overcoming Probabilistic AIs with Evidential Probability Relativism

By inflating probabilistic truth in the manner described above, we are able to derive a solution to probabilistic AIs that does not face the asymmetry challenge (i.e. the challenge of explaining how nominally probabilistic alternatives could be less relevant relative to thoroughly probabilistic beliefs, when they are clearly relevant relative to the corresponding nominally probabilistic beliefs held at the same context).

EPR can circumvent this challenge by admitting that the nominally probabilistic content of Thilo's ticket (certainly) winning *is* a relevant alternative to Thilo's thoroughly probabilistic belief. We can do so without having to accept skepticism, because EPR's inflated account of probabilistic truth allows us to deny AI2, the claim that Thilo cannot rule out that he has (certainly) won.

In assessing whether Thilo is in a position to rule out the alternative of winning, two factors are relevant. First, we need to consider whether and in virtue of what Thilo's probabilistic belief counts as true or false. According to EPR this verdict is dependent on our knowledge about the case, since this is what determines the correct evidential probability reference point. Secondly, we have to consider whether Thilo's evidence is sufficient to lead him to the correct conclusion.

Given the initial description of Lottery, for all we know there are 999 losers and one winner. Our total evidence determines a probability space,  $\psi$ , according to which the proposition that Thilo lost is assigned a value of .999. According to  $\psi$  it is true that Thilo's ticket probably lost ( $\lambda$ ) and false that it has (certainly) won ( $W$ ).<sup>31</sup> Given that, in this case, it is the evidential probabilities determined by our total evidence that constitute the measure of truth, our merely statistical evidence suffices for us to see that it is not the case that Thilo's ticket has (certainly) won. Believing that Thilo has (certainly) won is not consistent with the credences an ideally rational agent would have. It is thereby false.<sup>32</sup> Since we know that Thilo has the same evidence regarding the outcome of the lottery as we do, we can conclude that he is in a position to rule out the challenging nominally probabilistic alternative of winning as well.

Were we to know that Thilo's ticket has in fact won, we would have to conclude that Thilo is not in a position to rule out this alternative. This is in line with ACT. Still, we could recognize that *relative* to Thilo's perspective things look different. Relativism allows for us to pay due respect to multiple points of view.<sup>33</sup> Though we have to judge that Thilo cannot rule out the alternative of winning relative to our perspective, because it is true that he has (certainly) won, we can still recognize that relative to his perspective ruling out the possibility of having (certainly) won is in some sense adequate.<sup>34</sup>

<sup>31</sup>Because  $\psi \subseteq \lambda$ , but  $\psi \not\subseteq W$ .

<sup>32</sup>It is important to note that this does not imply that we are in a position to know that Thilo's ticket (certainly) lost. This nominally probabilistic content is equally false relative to  $\psi$ . This fits well with the earlier observation that we cannot know a nominally probabilistic content on the basis of merely statistical evidence; even if the evidence implies that the content is very likely true. Additionally, one might worry that from what I say it follows that relative to  $\psi$  we are in a position to know that Thilo has neither lost nor won. Admittedly, this would be a bad result. However, it does not follow. Relative to  $\psi$  we are in a position to know that 'it is not the case that Thilo *certainly* lost or *certainly* won'. This negated disjunction of nominally probabilistic contents is not equivalent to the type-shifted version of the propositional compound 'it is not the case that Thilo has lost or won,' for the type-shifted version of this content would read as 'it is *certainly* not the case that Thilo lost or won.' See Moss (2018b: §3.5) for details.

<sup>33</sup>See MacFarlane (2014: Ch. 5, §12.1).

<sup>34</sup>What if we, as opposed to Thilo, knew that he (certainly) lost? We would have to judge that Thilo's belief is trivially true. Does this also imply that we should judge that Thilo is in a position to rule out

To sum up, by inflating our account of probabilistic truth in the above specified way, we are able to overcome the challenge from probabilistic AIs as well as the asymmetry challenge. As set out in §2, it is hard to find compelling reasons supporting GAR. The proponent of EPR for probabilistic contents concedes that GAR is false. Instead she argues that asymmetry is to be found on the truth-conditional level. Propositional contents are true when they match the world in which (or about which) they are uttered. When Thilo believes that his ticket (certainly) lost, this belief will be true in virtue of matching the way the world really is. Merely statistical evidence does not suffice for Thilo to decide whether he is in a winning or losing world. It does, however, suffice for Thilo to recognize that the evidential probabilities are not extreme.

#### 4. Evidential Probability Relativism and Probabilistic Gettier Cases

We might not want to stop here. Though EPR as an inflated account of probabilistic truth does well in defying probabilistic AIs and the asymmetry challenge, it blocks us from a natural assessment of probabilistic Gettier situations. To see this consider the following case:

##### **Fake Lottery Letters**

Alice receives a letter from her town council informing her that each member of town is participating in a free lottery. The letter specifies that whether a citizen is a winner or a loser has been decided by the following procedure: for each citizen a random picking mechanism drew one ball from an urn containing 999 balls marked as 'loser' and one marked as 'winner'. The letter further states that winners will be announced at the next town meeting. Based on the information provided in the letter, Alice comes to believe that she probably lost. However, Alice does not know that her letter – and all other letters sent out by town officials – intentionally contain false information. To give people a nice surprise, town officials in fact held an easy lottery including 999 balls marked as 'winner' and only one ball marked as 'loser'. Alice's belief nonetheless turns out to be true. Some meticulous secretaries, angered by the concept of an easy lottery, were able to secretly change the ratio of winning to losing balls to 1 to 999 for a small number of draws including the one that decided whether Alice was a winner or a loser.

Here is a dilemma for the relativist who picks evidential probabilities as reference points. It seems that Alice does not know that she probably lost, because it appears to be a matter of pure luck that Alice's lottery draw was indeed standard and not easy to win. If the latter would have been the case, intuitively, Alice's probabilistic belief would have been misled. The only way a proponent of EPR can account for this verdict, as far as I see, is this: she can argue that while relative to Alice's perspective she ends up with a true belief in the standard as well as the easy lottery scenario simply by following her total evidence concerning the lottery, relative to our perspective and our superior knowledge of the situation Alice ends up with a false belief if she follows the evidence

---

the alternative of winning? It appears to me that this is not the case. We have to judge that Thilo's belief is merely true because it fits the way the world really is. The relevant reference point is a divine probability space (i.e. the analog of a possible world). Thilo's evidence, however, does not suffice for him to determine what kind of a world he is in: a losing or a winning world. It therefore does not suffice for him to rule out the alternative of winning. Still, we can acknowledge Thilo's perspective and make sense of why he holds the probabilistic belief nonetheless.

from the letter in an easy lottery situation. That is, relative to our perspective Alice receives *misleading evidence* in the easy lottery scenario, where misleading evidence is defined as evidence that leads her to believe something false. Since, relative to our perspective, it appears to be a matter of pure luck that Alice was not misled; her true, justified probabilistic belief is Gettiered.

While this at first might seem like a good explanation of why we hold that Alice's belief does not constitute knowledge, it ultimately leads us to problematic conclusions concerning what counts as misleading evidence in connection with probabilistic beliefs. Revisit, once again, Lottery. Imagine that we already know that Thilo is the lucky winner, while Thilo is oblivious to this fact. For all he knows his ticket probably lost. Given the above explanation of why it is true that Alice could easily have been misled, we would have to conclude that Thilo in believing that his ticket probably lost *is* misled by his evidence. After all, he is led to believe something that is false relative to our perspective. But this seems wrong. Thilo has only correct, unambiguous information about the constitution and procedure of the lottery. I hold that what we would like to conclude in his case is that, though his belief is false, the evidence upon which he formed his high credence did not lead him astray either. Again, truth-relativism allows us to consider multiple perspectives. For some questions – like truth *simpliciter* – the assessor's perspective will be decisive, while for other questions – like whether a subject was misled by her own evidence – the perspective from the context of use will be key.

If we accept that Thilo is not misled by his statistical evidence, because relative to his perspective he, epistemically speaking, ends up in the correct place, the same has to hold for Alice. In the easy lottery scenario Alice is still receiving a letter which indicates that she participated in a standard lottery. Though this information is false in this case, relative to Alice's perspective her probabilistic beliefs are not led astray when they conform with her evidential probabilities. And her evidential probabilities indeed indicate that she has probably lost. The dilemma for the proponent of EPR now takes the following form: either she concludes that just as Alice is misled by the information in the letter in the easy lottery scenario, Thilo is misled by the statistical evidence when he has in fact won the lottery, or she concedes that neither Alice nor Thilo are misled in these cases. Both answers are not satisfactory.

The worry for EPR, then, is the following. Regarding the standard lottery case, it seems plausible that although we would consider Thilo's probabilistic belief to be false if he actually won, we would not go so far as to judge that Thilo was misled by the statistical evidence, since the lottery was indeed designed to make it very unlikely that he would win. This suggests that in assessing whether our probabilistic beliefs are led astray by our total evidence, our perspective prevails. EPR, however, does not leave any room for such a view of being misled by one's evidence. Under EPR, (rationally) following our total evidence can never lead us astray, no matter whether our evidence consists in obviously misleading information (as for Alice in the easy lottery scenario) or not (as for Thilo in the standard lottery case). EPR thereby blocks us from the intuitive explanation that Alice's probabilistic belief is Gettiered *because* her belief could easily have been misled.

The good news is that the problem does not lie within relativism itself, but only with EPR's choice of the type of reference points. We do not think that Thilo gets something right merely because he correctly follows his evidence. The same would be true for Alice, but in her case we believe that she easily could have been misled. The important difference between Thilo and Alice is that Thilo's statistical evidence properly functions as an indicator of the actual objective chances of his ticket losing before the draw, while

Alice's evidence in the easy lottery scenario leads her away from what is in fact objectively likely. To accommodate for this difference, we need to provide a theory of probabilistic truth that centers on objective chances rather than evidential probabilities as reference points.

## 5. The Objective Chance Analog

A simple suggestion on how to conceive of probabilistic truth in terms of objective chance is given by Hájek (2011) who presents the following analogy (from now on referred to as the objective chance analog or OCA): “[t]ruth is to belief, as agreement with objective chance is to degree of belief.” The claim is intuitively appealing. As Hájek points out, full beliefs are quite obviously governed by a norm of veracity. Imagine two agents, each entertaining a consistent set of beliefs. While one agent entertains mainly false beliefs, the other has mainly true ones. A natural verdict is that the latter agent is doing better, epistemically speaking. Something similar applies to our credences. For instance, we think that some weather forecasters are better than others without putting into question that the probabilistic beliefs that they communicate are coherent. Furthermore, adopting Hájek's simple approach puts us in a position to make out an important difference between Lottery and Fake Lottery Letters. As stated earlier, Thilo's evidence is not misleading concerning the objective chances of his ticket being a loser, while Alice's evidence could have easily indicated something false concerning the objective chances of her being a loser in the communal lottery. There are, however, obstacles to this approach.

The first arises from the fact that OCA in its rather loose formulation does not tell us enough about *which* objective chances are relevant for determining probabilistic truth. A natural clarification to the demand is the following: Credences (and other probabilistic beliefs) possess the probabilistic analog of truth iff they match (or are consistent with) the objective chances *at the world and time which the belief in question is about*. This response yields two major problems.

First, if probabilistic beliefs are true *at a world* in virtue of matching certain objective chances at this world, it is not clear what makes them different from simple propositional beliefs *about objective chances*. However, as Moss (2018b: 2) points out, probabilistic beliefs “are not full beliefs about objective chance facts.” Moss shows that in using probabilistic or epistemic vocabulary we frequently do express something different than mere beliefs about chances. The naive version of the objective chance analog is, in this sense, a variety of contextualism. It, therefore, falls prey to the same general objection against contextualist theories of probabilistic truth discussed in §3.1; it inadvisably conflates the two notions of probabilistic and full beliefs, leaving no room for any substantial differences.

Secondly, objective chances of past events are often assumed to be extreme. From this it follows that our probabilistic beliefs about the past could only ever be trivially true, while our moderate credences about the past could never be true at all. This is false. If credences are ever fit for constituting knowledge, Thilo's rational .999 credence that his ticket *has lost* should be a clear instance of a knowledge-apt probabilistic belief. But in believing that it is .999 likely that his ticket has lost, Thilo entertains a probabilistic belief about a past event. The draw already took place and his ticket was either chosen to be the winner or not.<sup>35</sup> Nonetheless, Thilo's credence appears to be correct. A

<sup>35</sup>See Moss (2018b: 129–30) and Pettigrew (2017: 107, 116–21).

proper theory of probabilistic truth should allow for this correctness to reflect in the veracity of Thilo's belief.

In the next section I present a version of the objective chance analog for which both these issues do not arise. My proposal is simple. We should adopt a variant of truth-relativism which centers on objective chance rather than evidential probabilities. I call this kind of relativism *objective chance relativism* (OCR).

### 5.1. Objective Chance and Relative Truth

Remember that one of our initial observations was the following: Probabilistic claims are claims that include epistemic vocabulary, and epistemic vocabulary is sensitive to a subject's epistemic state or evidence. We further saw that it was advisable to determine the truth of probabilistic contents from the context of assessment rather than from the context of use. So far we have only considered the possibility of an assessor's, *A*'s, total evidence determining a reference point that mirrors *A*'s evidential probabilities. This is not the only viable option we have, though it might be the most obvious. I argue that objective chance reference points, similar to evidential probability reference points, are sensitive to *A*'s total evidence in such a way that a change in *A*'s total evidence potentially invokes shifts from one objective chance reference point to another.

To arrive at this conclusion it is important to notice that objective chances are indexed to worlds and times. An event can have multiple objective chances at different worlds and, most importantly, different times. Intuitively, in Lottery Thilo can know that his ticket has probably lost even after the draw has already taken place. Though the objective chances of Thilo having lost are either 1 or 0 after the draw, Thilo's belief appears to be non-trivially true. One immediate conjecture as to how this result could be realized on a version of OCA is to postulate that objective chance reference points are in some way connected to our (the assessor's) total evidence about Thilo's case, such that, given our total evidence, the appropriate reference point is the reference point that mirrors the non-extreme objective chances of Thilo's ticket losing *shortly before* the draw.

The upshot is that postulating a connection between a subject's total evidence concerning a certain event and the objective chance of the event's occurrence *at a certain time* is in no way *ad hoc*. To the contrary, it can be directly read off of Lewis's (1987: 87) independently plausible Principal Principle:

#### Principal Principle (PP)

Let  $P_0$  be the credence function of an agent at the beginning of her epistemic life. Let  $t$  be any time. Let  $\omega$  be any world. Let  $x$  be a real number in the unit interval. Let  $C_{t, \omega}(A) = x$  be the proposition that the chance at time  $t$  and world  $\omega$  of  $A$ 's holding equals  $x$ . Let  $E$  be the agent's total evidence that is admissible at time  $t$ . Then

$$P_0(A | C_{t, \omega}(A) = x \ \& \ E) = x$$

PP relates time-dependent objective chance to time-dependent admissible evidence via the concept of rational credence.<sup>36</sup> Admissible evidence, in turn, is defined as follows:

<sup>36</sup>See Lewis (1987: 114).



Evidence  $E$  is admissible at  $t$  concerning outcome  $A$  iff it contains the sort of information whose impact on credence about  $A$ , if any, comes entirely by way of credence about the chances of  $A$  at  $t$ .<sup>37</sup>

Concerning Lottery, let  $L$  be the proposition that Thilo's ticket loses. Let  $t_{-1}$  be the time shortly before the draw and  $C_{t_{-1}, \omega}(L) = .999$  the proposition that at world  $\omega$  and time  $t_{-1}$  the objective chances of Thilo's ticket losing are .999. Let  $E$  be our total evidence. We can now investigate the following question: Which credence  $P_0$  in  $L$  should we adopt at the present instance  $t_0$ ? Following PP we should reason as follows:

(P1) If our present total evidence  $E$  is admissible at  $t_{-1}$  concerning  $L$ , then the present rational credence for us is  $P_0(L \mid C_{t_{-1}, \omega}(L) = .999 \ \& \ E) = .999$ . (according to PP)

(P2)  $E$  is admissible at  $t_{-1}$  concerning  $L$ .

(C) Therefore, the present rational credence for us is  $P_0(L \mid C_{t_{-1}, \omega}(L) = .999 \ \& \ E) = .999$ .<sup>38</sup>

$E$  is admissible at  $t_{-1}$  because it contains only information about the objective chances of  $L$  at  $t_{-1}$ .  $E$  does not include any information concerning the specific outcome of the draw. If we had such information,  $E$  would no longer be admissible at  $t_{-1}$ , but only at a later point in time *after* the draw.

Together with PP it follows that it is rational for us to have a .999 credence that Thilo lost, though the draw already took place. The relativized version of the objective chance analog (OCR) simply adds that we should adjust our credences in accordance with PP, because given that we are not misled about the objective chances of Thilo losing being .999 at  $t_{-1}$ , adjusting our credence in this way will satisfy the probabilistic analog of truth (at least from our perspective). Consequently, we are able to judge that Thilo's probabilistic belief is non-trivially true. It is consistent with the objective chance reference point determined by our total evidence concerning whether  $L$ .

My suggestions, therefore, amount to the subsequent relativistic principle for determining objective chance reference points: objective chance reference points shift relative to the admissibility of a subject's total evidence concerning a certain question at the context of assessment of a probabilistic belief, assertion, or content. One might worry that a subject's total evidence is admissible at multiple points in time concerning any outcome  $A$ . For instance, our total evidence  $E$  is admissible at  $t_{-1}$  concerning  $L$ . But according to the above definition of admissibility it is also admissible at later times, specifically at the present time  $t_0$  after the draw. What is it then that connects  $E$  to the objective chances at  $t_{-1}$  rather than  $t_0$ ? As a rule, a subject's total evidence with regard to a certain outcome is admissible during some interval  $[t_e, \infty)$  that begins at some time  $t_e$  and extends continuously into the future. This implies that for each outcome  $A$  there is a certain point in time where a subject's total evidence *starts* being admissible. Adopting Lewis's terminology, call the earliest point in the interval the *endpoint time of admissibility* of  $S$ 's total evidence in relation to  $A$ . It is this endpoint time of admissibility that determines the correct objective chance reference point. We end up with a probability space mirroring the chances of  $A$  at  $t_e$ .

<sup>37</sup>Derived from Lewis (1987: 92) and Hoefler (2007: 553).

<sup>38</sup>See Lewis (1987: 114–16) for a similar application of PP.

### 5.2. Advantages of Objective Chance Relativism

With OCR in place, we are in a position to review multiple advantages of the view. First and foremost, OCR just like EPR is fit for overcoming the challenge from probabilistic AIs without introducing the asymmetry challenge. Our total information concerning the outcome of the draw has its endpoint time of admissibility before the draw. Consequently, our total information determines an objective chance reference point that mirrors the objective chances of Thilo losing shortly before the draw (i.e. .999). Relative to this reference point it is clearly false that Thilo has (certainly) won. Our merely statistical evidence suffices for us to see this. There are 1000 tickets in total, one winning ticket, and a randomized picking mechanism. Thilo's ticket is not a certain winner. It is probably a loser. Since we know that Thilo has the exact same evidence concerning the draw as we do, we can conclude that he is in a position to rule out the alternative of (certainly) winning as well.

Additionally, in centering our variant of relativism on objective chances, we are in a position to make sense of probabilistic Gettier cases. We can now explain our intuition that in Fake Lottery Letters Alice could have easily been misled by her evidence without having to deny that in Lottery Thilo is not misled by the statistical evidence, even if he has won. If Alice's lottery draw had been easy (as intended by the town officials) she would have believed that she probably lost even though the objective chances determined by the (endpoint time of) admissibility of her total evidence, regarding whether she lost would have been very low (namely .001). Relative to her perspective, Alice, epistemically speaking, would not have ended up where she intended to. This is different for Thilo. Even in the case where he wins the lottery (and we know this) it is false that Thilo was misled by his evidence. As argued above, in investigating whether a subject has been misled by her evidence, her perspective prevails. Thilo has only information about a time shortly before the draw took place. At this time the objective chances of losing were indeed very high (i.e. .999). This is true even for the case where, at the present, Thilo is the lucky winner. So according to Thilo's perspective, given that he does not have any information about the outcome of the draw, even if he has in fact won, his high credence is not misled (though ultimately false). He would have been led astray, however, if the objective chances of losing *shortly before the draw* had been different than indicated by his evidence. This would be the case, for instance, had Thilo received misleading information concerning the total number of tickets or the unbiasedness of the picking mechanism.<sup>39</sup>

---

<sup>39</sup>It is important to note that OCR has an advantage over EPR in this case only if we intend to analyze probabilistic Gettier situations using the concept of misleading evidence. I have argued that one intuitive way to identify Alice's belief in Fake Lottery Letters as Gettierized is via the judgment that she could have easily been misled by the false information provided by the letter and that EPR blocks us from reaching this verdict. A proponent of EPR, however, still has the opportunity to present a way to identify Gettierized probabilistic beliefs without making use of the concept of misleading evidence, and maybe even without having to refer to probabilistic truth altogether. Thanks to an anonymous referee for pointing this out to me. Even if the proponent of EPR can sidestep the need to talk about misleading evidence in connection with probabilistic beliefs, I believe that proposing OCR as a genuine rival account to EPR is still fruitful. OCR has further advantages over EPR besides handling Gettier situations in a certain way (see below). On the other hand, there might also turn out to be areas and issues where EPR fares better than OCR. An extensive discussion of all the advantages and disadvantages of EPR and OCR exceeds the scope of this paper. However, with both versions of probabilistic truth in place, such a discussion can be considered a task for further research.

Finally, third-person assessments, retractions and disputes are still explainable with OCR. If we receive evidence about the outcome of the lottery, by reading the newspaper for example, our objective chance reference point shifts such that it now mirrors the extreme objective chances of the ticket winning *after* the draw. Imagine we learn from the newspaper that Thilo has won. Let us further assume that this is indeed true. Our reference point is now a divine probability space containing only the actual world where Thilo has won. If Thilo has not gained any further evidence his reference point and ours will again diverge. Since truth is determined at the context of assessment we are able to conclude (as desired) that Thilo's probabilistic belief is false.<sup>40</sup>

At the same time OCR does not relativize too much. On the view it is easily possible that multiple bodies of total evidence refer to the same objective chance reference point. There are, thereby, many situations where subjects with diverging bodies of evidence nonetheless share a reference point with regard to a certain probabilistic content. This allows us to reintroduce a kind of unrelativized objectivity that Moss appears to be sympathetic to. Consider one last case discussed in Moss (2018b: 21):

### Criminals in Paris

Some spies in London are eavesdropping on some criminals in Paris. The criminals are trying to figure out where James Bond is located. The eavesdropping spies are confident that Bond is in London. They overhear that the criminals have collected a lot of misleading evidence suggesting that Bond is in Paris. One criminal asserts: "It is unlikely that Bond is in London." The spies immediately object: "That's false. Bond is almost certainly in London."

Concerning this case Moss (2018b: 128) emphasizes that "[t]he criminal may be justified in believing the content that Bond is probably in London. But the content of her belief is not true for her, or indeed true in any sense."

The upshot is that OCR can reconcile Moss's verdict about the case with a relativistic approach on probabilistic truth. Given the case description, it is plausible that the criminal's and the eavesdropper's total bodies of evidence do have the same endpoint time of admissibility,  $t_e$ , concerning the question of Bond's whereabouts. Therefore, their probabilistic utterances and beliefs concerning Bond's location will both be true or false relative to *the same* reference point. Given that the objective chances of Bond being in London at  $t_e$  are actually very high, or even 1, it is true that, in this situation, the criminal's utterance does not count as true in any sense. It is false from her and from the eavesdropper's perspective.

A note of caution is in order. Until now I have fully relied on the reader's intuitions concerning whether and which objective chances are present in Lottery and other cases. Completing the here presented account of OCR with a concrete analysis of objective chance remains an open task. For now, I can say this much. There appears to be a rather widespread agreement that non-extreme objective chances exist. Further, many authors hold that there are non-extreme objective chances for a wide variety of everyday phenomena.<sup>41</sup> No matter how broad a conception of objective chances will be, it might

<sup>40</sup>Note that if our newspaper evidence is misleading such that it indicates that Thilo has lost, our reference point nonetheless remains unchanged. We are therefore misled and our belief that Thilo lost, as well as our potential judgement that Thilo's probabilistic belief is trivially true, are false. I take this to be a desirable result.

<sup>41</sup>See, for instance, Strevens (1999) and Hoefer (2007), who defend the presence of objective chance in a variety of fields. There appear to be objective chances in lotteries, card games, and for gambling devices, but

occur that not all (non-past) events do have non-extreme objective chances. As a direct consequence OCR predicts that there are certain probabilistic beliefs that will only ever be trivially true and certain moderate credences that will never be true. I do not think that this result is problematic. The claim that I could know that it is *exactly* .6 likely that Jones smokes appears rather puzzling once it turns out that there is no moderate probability or chance that matches the content I claim to know. How else could it be proper for me to claim that I know something this specific? If there is no moderate objective chance reference point, because the objective chance of Jones smoking is 1 for instance, then OCR still allows me to come to know less precise probabilistic contents. I can, for example, come to know that Jones probably smokes, that it is more likely than not that Jones smokes, or even that it is *at least* .6 likely that she does.

## 6. Conclusion

Probabilistic knowledge is an intriguing concept that has the potential to find wide application in epistemology. The aim of this paper is to provide a theory of probabilistic truth that supplements Moss's theory of probabilistic knowledge in a way that allows us to overcome challenges that the initial deflationary approach on probabilistic truth cannot adequately conquer. The hope is that overcoming these challenges will make probabilistic knowledge more palatable to epistemologists. As pointed out, deflationism not only leaves us ill-equipped to overcome probabilistic AIs, it also is perceived as a rather opaque answer to the factivity challenge. OCR thus displays two virtues: it allows us to give a straightforward answer to probabilistic AIs without forcing us to introduce and explain a dubious asymmetry between thoroughly and nominally probabilistic beliefs and, secondly, it provides us with a clear sense of when and in virtue of what probabilistic contents are non-trivially true. Additionally, OCR, as opposed to its rival EPR, allows us to assess probabilistic Gettier cases in an intuitive way and leaves room for the judgment that in certain situations probabilistic beliefs, even when rational or justified, might not be true in any sense. On the here presented approach, probabilistic truth can only be as graspable as our concept of objective chance is. Although I am confident that a variety of phenomena have objective chances, further investigation is needed to make probabilistic truth, and consequently probabilistic knowledge, easily manageable.

## References

- Albert D.Z. (1992). *Quantum Mechanics and Experience*. Cambridge, MA: Harvard University Press.
- Blackburn S. (1984). *Spreading the Word: Groundings in Philosophy of Language*. Oxford: Oxford University Press.
- Demey L., Kooi B. and Sack J. (2019). 'Logic and Probability.' In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2019 edition). <https://plato.stanford.edu/archives/sum2019/entries/logic-probability/>
- DeRose K. (1995). 'Solving the Skeptical Problem.' *Philosophical Review* 104(1), 1–52.

---

also in quantum mechanics, thermodynamics, chemistry, and medicine. Hofer (2007: 557) goes as far as to suggest that we might be able to assign my success of catching the 9:37 train on a weekday a non-extreme objective chance. Additionally, by Hofer's account objective chances exist irrespective of whether determinism is true or not. See also Albert (1992) and Loewer (2020) who defend "the Mentaculus" framework according to which it is possible to determine the objective chances of A given B for *all* physically specifiable propositions A and B.

- Easwaran K.** (2018). 'Review of the Book: Probabilistic Knowledge, by S. Moss.' *Notre Dame Philosophical Review*. <https://ndpr.nd.edu/reviews/probabilistic-knowledge/>
- Eder A.-M.A.** (2019). 'Evidential Probabilities and Credences.' *British Journal of the Philosophy of Science*, 1–21. <https://doi.org/10.1093/bjps/axz043>
- Egan A.** (2007). 'Epistemic Modals, Relativism and Assertion.' *Philosophical Studies* 133(1), 1–22.
- Gibbard A.** (2003). *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Goldman A.I.** (1988). 'Psychology and Philosophical Analysis.' *Proceedings of the Aristotelian Society* 89, 195–209.
- Hájek A.** (2011). 'A Puzzle about Degree of Belief.' Unpublished Manuscript.
- Hawthorne J.** (2005). 'The Case for Closure.' In M. Steup, J. Turri and E. Sosa (eds), *Contemporary Debates in Epistemology*, pp. 26–34. Chichester: Wiley-Blackwell.
- Hoefer C.** (2007). 'The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance.' *Mind* 116, 549–96.
- Kratzer A.** (1991). 'Modality.' In A. v. Stechow and D. Wunderlich (eds), *Semantics: An International Handbook of Contemporary Research*, pp. 639–50. Berlin: de Gruyter.
- Kyburg H.E.** (1961). *Probability and the Logic of Rational Belief*. Middletown, CT: Wesleyan University Press.
- Lewis D.** (ed.) (1987). 'A Subjectivist's Guide to Objective Chance.' In *Philosophical Papers Volume II*, pp. 83–131. New York, NY: Oxford University Press.
- Lewis D.** (1996). 'Elusive Knowledge.' *Australasian Journal of Philosophy* 74(4), 549–67.
- Loewer B.** (2020). 'The Mentaculus Vision.' In V. Allori (ed.), *Statistical Mechanics and Scientific Explanation: Determinism, Indeterminism and Laws of Nature*, pp. 3–29. Singapore: World Scientific.
- MacFarlane J.** (2011). 'Epistemic Modals are Assessment-Sensitive.' In A. Egan and B. Weatherson (eds), *Epistemic Modality*, pp. 144–78. Oxford: Oxford University Press.
- MacFarlane J.** (2014). *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.
- MacFarlane J.** (2019). 'On Probabilistic Knowledge.' *Res Philosophica* 1, 97–108.
- Moss S.** (2018a). 'IX – Moral Encroachment.' *Proceedings of the Aristotelian Society* 118(2), 177–205.
- Moss S.** (2018b). *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Moss S.** (2019). 'Full Belief and Loose Speech.' *Philosophical Review* 128(3), 255–91.
- Moss S.** (2021). 'Knowledge and Legal Proof.' In T. Szabó Gendler, J. Hawthorne and J. Chung (eds), *Oxford Studies in Epistemology Volume 7*, pp. 176–213. Oxford: Oxford University Press.
- Nelkin D.** (2000). 'The Lottery Paradox, Knowledge, and Rationality.' *Philosophical Review* 109(3), 373–409.
- Pettigrew R.** (2017). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Pritchard D.** (2008). 'Knowledge, Luck, and Lotteries.' In V. Hendricks and D. Pritchard (eds), *New Waves in Epistemology*, pp. 28–51. London: Palgrave Macmillan.
- Rich P.** (2020). 'The Logic of Probabilistic Knowledge.' *Philosophical Studies* 117, 1703–25.
- Smartt T.** (2019). 'From Outside of Ethics: Moss, Sarah. Probabilistic Knowledge.' *Ethics: An International Journal of Social, Political, and Legal Philosophy* 129(2), 430–8.
- Stephenson T.** (2007). 'Judge Dependence, Epistemic Modals, and Predicates of Personal Taste.' *Linguistics and Philosophy* 30(4), 487–525. <http://dx.doi.org/10.1007/s10988-008-9023-4>.
- Strevens M.** (1999). 'Objective Probability as a Guide to the World.' *Philosophical Studies* 95(3), 243–75.
- Turri J.** (2021). 'Knowledge Attributions and Lottery Cases: A Review and New Evidence.' In I. Douven (ed.), *Lotteries, Knowledge, and Rational Belief*, pp. 28–47. Cambridge: Cambridge University Press.
- Weatherson B.** (2011). 'Defending Interest-Relative Invariantism.' *Logos and Episteme* 2(4), 591–609.
- Williamson T.** (2002). *Knowledge and its Limits*. Oxford: Oxford University Press.

**Svenja Schimmelpfennig** is a graduate student at the University of Cologne in Germany where she earned her Bachelor's degree in Philosophy and Information Processing in 2020. She is currently working towards the completion of her Master's degree in Philosophy. Her philosophical interests lie primarily in the area of epistemology and in the overlapping areas of epistemology, philosophy of language, and metaphysics.

**Cite this article:** Schimmelpfennig S (2023). Probabilistic Truth, Relativism, and Objective Chance. *Episteme* 20, 757–777. <https://doi.org/10.1017/epi.2023.7>