



Infinite frequency principles of direct inference

Lennart B. Ackermans¹

Received: 24 June 2021 / Accepted: 1 March 2022
© The Author(s) 2022

Abstract

According to an infinite frequency principle, it is rational, under certain conditions, to set your credence in an outcome to the limiting frequency of that outcome if the experiment were repeated indefinitely. I argue that most infinite frequency principles are undesirable in at least one of the following ways: (1) accepting the principle would lead you to accept bets with sure losses, (2) the principle gives no guidance in the case of deterministic experiments like coin tosses and (3) the principle relies on a metaphysical property, ‘chanciness’, whose necessary and sufficient conditions are unknown. I show that a frequency principle that is based on the principal principle suffers from problems related to the definition of ‘chance’ or ‘chanciness’, which could lead to all three of the above problems. I introduce a version of the infinite frequency principle that does not rely on a notion of chance or chanciness and does not suffer from any of these problems.

Keywords Direct inference · Bayesianism · Subjective probability · Objective probability · Principal principle

1 Introduction

An infinite frequency principle prescribes that a rational agent should set her credence that an experiment has outcome a to the frequency of a 's if that experiment were repeated indefinitely, provided that she knows this frequency. Infinite frequency principles differ from each other concerning the conditions under which this holds. These principles play a prominent role in some philosophical theories of subjective probability, such as Howson and Urbach (2006) and Williamson (2010), as well as theories of objective probability (Mellor, 1995).

Infinite frequency principles can be used for what is known as *direct inference*, the calibration of one's credence using evidence of chances and frequencies. If all

✉ Lennart B. Ackermans
ackermans@esphil.eur.nl

¹ Erasmus School of Philosophy, Erasmus University Rotterdam, Rotterdam, The Netherlands

that rationality requires is adherence to the probability axioms, then there are no principles of direct inference. Although one *may* adhere to such a principle, one is not irrational otherwise. But rationality might require more. A typical example involves finite frequencies: suppose one is drawing from an urn with blue and red balls, with replacement. If all one knows is that the frequency of blue balls in the urn is q , it seems unreasonable to set one's credence in drawing a blue ball to anything else but q . A different kind of reasoning could be applied if one has knowledge about hypothetical frequencies. Suppose that all one knows is that the frequency of drawing a blue ball approaches q as more balls are drawn (with replacement). Then it seems unreasonable to set one's credence in drawing a blue ball to anything else but q . This type of reasoning draws on an infinite frequency principle.

Infinite frequency principles belong to a broader class of direct inference principles, which include the well-known principal principle and various actual frequency principles. Discussion in the literature tends to be focused on problems related to direct inference principles in general, to the principal principle, and to inference based on actual (rather than limiting) frequencies, whereas explicit discussion of infinite frequency principles is less common.¹ Nevertheless, infinite frequency principles are sometimes implicitly discussed in the context of the principal principle. While the principal principle connects credences with *chances*, it is sometimes argued for on the basis of an infinite frequency principle, which connects credences with *limiting frequencies* (see Strevens, 1999; Howson & Urbach, 2006, pp. 76–78; Mellor, 1995, pp. 44–49; Williamson, 2010, pp. 39–42; Albert, 2005). (If chances *are* limiting frequencies, then the principal principle is itself an infinite frequency principle, namely the Chancy Infinite Frequency Principle discussed in Sect. 5.)

Another way in which an infinite frequency principle could be used is to achieve a better understanding of the nature of chance, the nature of rationality, or the way in which both are connected. The connection between chances and credences is typically described by the principal principle, according to which you should set your credence to p given knowledge that the chance is p . The principal principle is often used to evaluate whether a particular theory of chance picks out the right chances—those that are capable of guiding our credences (some examples are Schaffer, 2003; Schaffer, 2007; Eagle, 2004; Hoefer, 2007). In order to use the principal principle in this way, one needs to make assumptions about rationality that give substance to the principle. A frequency principle could do this by giving an independent justification for calibrating one's credence to certain values, in situations in which one might also know the chance. Given that these chances must, in virtue of the principal principle, be equal to the calibrated credences, the values that the chances can take on are constrained by the frequency principle. This is described in more detail at the end of Sect. 2.3 and in the concluding Sect. 7.1.

However, the frequency principles advocated in the literature suffer from various problems. Some of them are false, and many proposed arguments for frequency prin-

¹ Recent discussions on direct inference principles in general have given much attention to problems related to the admissibility of additional evidence (Hawthorne et al., 2017; Landes et al., 2021; Wallmann & Hawthorne, 2020; Wallmann & Williamson, 2020; Fischer, 2021). There has also been discussion on the reference class problem in relation to direct inference (Thorn, 2012; Thorn, 2017; Wallmann, 2017).

ciples fail. This paper gives an overview of infinite frequency principles, how they can be argued for and what their respective merits are.

I introduce three desiderata of frequency principles. First, they should be *informative*, meaning that they constrain which credences are rational in a sufficiently large number of cases, including cases of deterministic experiments like coin tosses (the *informativeness desideratum*). Second, they should not constrain the credence functions we deem acceptable too much, such that credence functions that are actually rational are excluded or—worse—such that only irrational credence functions satisfy the frequency principle (the *rationality desideratum*). Third, they should be precise. Frequency principles should have an obvious, unambiguous interpretation (the *precision desideratum*).

I show that most frequency principles do not satisfy all three desiderata. The *Naive Infinite Frequency Principle (NIFP)* does not satisfy the rationality desideratum. The *Chancy Infinite Frequency Principle (CIFP)* does not satisfy at least one of the three desiderata, depending on how it is interpreted. I introduce the *Equal Credence Infinite Frequency Principle (EIFP)*, which I claim satisfies all three desiderata, and is more informative than the well-known *CIFP*.

I show that setting one's credence to the limiting frequency, given some experiment, is rational only if it is rational to assign identical credences to each individual run of the experiment. This means that a frequency principle must constrain the set of experiments to which it applies to those for which equal credences to repetitions are rational.

Section 2 introduces some important definitions, the desiderata of frequency principles and gives some examples of applications. Section 3 shows that the *NIFP* does not satisfy the rationality desideratum. Section 4 argues that the popular betting argument does not succeed in justifying frequency principles, even if previously expressed concerns by Strevens (1999) would be resolved. It also gives an alternative consistency argument. Section 5 introduces *CIFP* and examines its merits. Section 6 introduces and examines the merits of *EIFP*. Section 7 concludes and suggests how *EIFP* might give new insights into the nature of chance.

2 Preliminaries

2.1 Desiderata of frequency principles

Frequency principles can be part of a foundational theory of rational credences. Within such a theory, it is supposed that an agent's credences can be described by a credence function, but that only some of these functions are rational. Axioms and principles constrain the set of credence functions that can be rational. While Howson and Urbach (1993) argued that their infinite frequency principle (a version of the principal principle based on frequentist chance) was a direct corollary of a subjectivist framework—which accepts only consistency requirements of rationality—it is now generally agreed that such a principle constrains a rational agent's credences to a greater extent than consistency alone does (Albert 2005; Howson and Urbach 2006). My first desideratum of a frequency principle is that it is a principle of this type: it constrains the set of rational

credence functions beyond consistency. Moreover, the more a principle constrains, the more informative it is—which I take to be better, presuming the constraints are indeed rationality requirements (see also Williamson, 2010, pp. 72–74). This is the *informativeness desideratum*.²

An important feature of probabilities interpreted as rational credences is that they can be non-trivial—different from 0 or 1—even if they concern outcomes of deterministic experiments. I will say that an experiment described by F is deterministic if there could be a complete and correct description of the experiment at the time it is performed (but before it has revealed its outcome), $F \& G$, such that an $F \& G$ -experiment can have only one possible outcome. (G must describe the world as it is at the time the experiment is performed, so it cannot include a statement of the experiment's outcome.) Whether an objective chance of a deterministic experiment can be non-trivial is a matter of debate (see e.g., Schaffer, 2007; Glynn 2009; Pivato & List, 2015), but that debate doesn't extend to subjective probability. It is generally agreed that credences of deterministic experiments can be non-trivial when one's evidence is limited.

An example is a toss of a fair coin. Coin tosses are plausibly described (at least approximately) by a deterministic model (Suppes, 1987). In Suppes's model, one can use attributes of an individual coin toss—the coin's upward velocity and spin—to calculate with full accuracy whether it will land heads or tails. Hence, a coin toss could have no other outcome than its actual outcome given the full description of the toss. Deterministic situations like this one are those that credences deal with quite often, and frequency principles are helpful tools in guiding our credences in such situations. A frequency principle, in order to be sufficiently informative, should be able to guide our (non-trivial) credences in outcomes of such experiments.

Secondly, a frequency principle should constrain credence functions in the right way and it should not constrain too much. That is, it should not scratch functions that are in fact rational. It certainly should not scratch so many functions that all remaining functions are irrational. This is the *rationality desideratum*. There are multiple ways in which a frequency principle can be incorrect in this sense, but in this paper, especially the sure loss principle and related intuitions will be considered. The sure loss principle states that a credence function is irrational if it could lead the agent to accept a sequence of bets whose final result is a sure loss for the agent, as logically implied by her evidence. Hence, a frequency principle should not constrain the set of rational credence functions to such an extent that for each of the remaining functions, there is some (hypothetical) sequence of bets, which the agent would accept, leading to a sure loss. The sure loss principle is endorsed by a great deal of authors who rely on Dutch Book arguments (e.g., Ramsey, 1931; De Finetti, 1937; Howson & Urbach, 2006), although a number of problems have been raised (see Hájek, 2009a).

I also discuss examples and arguments in which there is something close to a sure loss, such as an almost sure loss (a very high chance of a loss) and a sure absence of a

² It should be noted that a principle can also be informative in a different sense, even if it does not constrain beyond consistency: it can helpfully inform us what previously accepted rationality constraints like consistency require, when this might not have been obvious otherwise. Since the frequency principles discussed in this paper do not differ significantly with respect to this latter sense of informativeness, it does not play a role in my arguments.

gain (an impossibility of a positive payoff) with a possibility of a loss. Such situations are less clearly irrational but nevertheless have at least some of the intuitive appeal of the sure loss principle.

Finally, a frequency principle should be as precise as possible. This is the *precision desideratum*.³ If the principle is to be of use within a larger theory of rational credences, it is best if it does not contain any ambiguities, such as when the concepts involved are open to widely diverging interpretations. A good example of what I mean with an ‘ambiguity’ can be found in the disagreement about whether there can be non-trivial chances of deterministic experiments. According to some, the toss of a fair coin has a chance of 1/2 of landing heads, while others contend that the chance can only be 0 or 1. This disagreement is possibly a result of differing notions of ‘chance’, and the conflict would be resolved if ‘chance’ is given a precise definition. Now consider what happens when the concept of ‘chance’ is invoked in a principle of rational credence—such as the principal principle, which connects credences with chances. Two people who disagree along the above lines about what the chances are will give quite different recommendations of how to set your credence on the basis of the principal principle. If we are to come up with workable procedures for assigning rational credences in practice, such disagreements need to be resolved. One can do so by resolving the conceptual ambiguity.

Precision could also be a desideratum if one wants to use a frequency principle to better understand the nature of chance and its connection to credences. As I describe below, a frequency principle could be used in conjunction with the principal principle to test whether a theory of chance is correct (Sect. 2.3), or to give a definition of chanciness (Sect. 7.1). When the concepts involved are open to many conflicting interpretations, we cannot achieve a great deal of understanding into the nature of chance.

Not everyone may consider the precision desideratum to be important, but theorists who rely on frequency principles typically do. It is, for example, expressed by Howson and Urbach when they introduce their frequency principle:

There is a good deal of evidence that in suitable experimental contexts the relative frequency with which each of the various possible outcomes occurs settles down within a smaller and smaller characteristic interval as the number of observations increases. This is not very precise, however (how should the interval vary as a function of sample size?), and we shall follow [Mises (1957)] in replacing the rather vague notion of relative frequency “settling down within an increasingly small interval” by the precise mathematical definition of a limit. (Howson & Urbach, 2006, p. 46)

The precise definition of chance inspired by von Mises then allows Howson and Urbach to formulate a precise version of the Chancy Infinite Frequency Principle (see Sect. 5).

³ Note that I use the word *precision* in a different sense than advocates of imprecise probabilities, which are probabilities with interval rather than point values. Imprecise probabilities can still be precise in my sense of the term.

2.2 Credences and frequencies

Let $C(a_i, \text{ given } E)$ denote a credence function. This credence function describes an agent's credences, as a number between 0 and 1, that a unique event a_i occurs, supposing that the agent's evidence is E (her set of beliefs with degree 1). To avoid confusion, I reserve the notation $C(a_i | x)$ for the conditional probability of a_i given x , which could be interpreted in a variety of ways, such as the agent's betting quotient for a conditional bet that is cancelled if x is false, or just as the evidence-informed probability $C(a_i, \text{ given } x)$. However, evidence-informed and conditional probability need not overlap.⁴

Unique events, such as *the coin I am about to toss lands heads* are referred to by indexed lowercase letters like a_1 . An associated outcome type, such as *a coin lands heads* is referred to by the same lowercase letter without an index, such as a . Descriptions of an experimental setup, such as *a coin is tossed* are given uppercase letters such as F . The evidence that the experiment preceding a_i occurs and satisfies F is denoted F_i . When an experiment F , with possible outcome a , is repeated, each outcome event is indexed by the number of repetitions. For example, a_5 is the event that the fifth run of an experiment F leads to outcome a .

I assume that there is a relation between a rational agent's credences and the way she behaves in betting games. Suppose that an agent is participating in a betting game in which she can choose to buy or sell bets. The price, as well as the payoff of these bets, is denoted in utility (rather than money), from the perspective of the agent. If the agent's credence is equal to $C(a_1, \text{ given } E) = q$, then she would both buy and sell a bet for a price of qS that pays out S (an amount in utility) if a_1 occurs. Here q is also called the *betting quotient*. In other words, let $I(a_1)$ be the indicator function that returns 1 if a_1 occurs and 0 if a_1 does not occur. Then the agent would accept bets whose actual value (the total amount of utility gained or lost) is $S(I(a_1) - q)$, where the stake S can be positive or negative. I do not suppose that this condition of credences is universal—there might be cases in which it is irrational to participate in betting games at all—but I assume that *if* the agent is participating in a betting game, then her credences tell her which bets to accept.

A frequency principle prescribes that an agent who has information about limiting frequencies should use that information to calibrate her credences. Given a repeatable experimental setup F which describes an experiment with a possible outcome a , the hypothetical limiting frequency of outcome a when F is repeated indefinitely is denoted $f_\infty^F(a)$. Let E be any other "admissible" evidence. A frequency principle prescribes that under conditions Q , the agent's credence function should satisfy, for all $i > 0$,

$$C(a_i, \text{ given } F_i \text{ and } f_\infty^F(a) = q \text{ and } E) = q.$$

Different frequency principles state different conditions Q . What the proper conditions are will be investigated in the subsequent sections.

⁴ Conditional credence should satisfy $C(a \text{ and } b) = C(a | b)C(b)$. In objective Bayesianism, there is no such requirement for $C(a_i, \text{ given } E)$.

Admissible evidence is the type of evidence that rationality would allow you to disregard when F_i and $f_\infty^F(a) = q$ are also part of your evidence; that is, F_i and $f_\infty^F(a) = q$ “trump” E with respect to a_1 . Note that this definition of admissibility is different from Lewis’ definition of admissibility with respect to the principal principle.⁵ In the case that a frequency principle is restricted to chance experiments F (see the **CIFP** in Sect. 5), one might instead use Lewis’ definition of admissibility. This paper will not deal with the question of how admissibility is best defined.

Although a version of the principal principle can be equivalent to the Chancy Infinite Frequency Principle (Sect. 5), the principal principle is not a frequency principle. The original principal principle proposed by Lewis (1980) refers to single case chances, which are only indirectly related to frequencies. The version of the principal principle (**PP**) I discuss in this paper instead uses a type chance $ch_F(a)$, which denotes the chance that an experiment described by F yields outcome type a .

PP: C is rational only if for all F , a and admissible evidence E we have $C(a_i, \text{ given } F_i, ch_F(a) = p \text{ and } E) = p$.

Before we can discuss these frequency principles, two interpretative questions must be addressed. First, what does it mean to repeat an experiment? The evidence F describes an experiment like a coin toss. An experiment F is repeatable if it is possible for an agent to be in consecutive situations in which her evidence with respect to the outcome a_i is given by F_i each time. However, there are multiple ways in which an experiment can be repeated, and these might lead to different limiting frequencies.

For example, suppose that you toss a fair coin in a normal way, described by F . As is commonly accepted, the limiting frequency of ‘heads’ of tossing fair coins is $1/2$. But now suppose that a robot is sitting next to you and calculates, whenever the coin is about to leave your hand, whether it will land heads or tails. If it is going to land tails, its robotic arms stop you from tossing the coin. With this way of repeating F , the limiting frequency of ‘heads’ is 1. We can’t have both $f_\infty^F(\text{heads}) = 1/2$ and $f_\infty^F(\text{heads}) = 1$, so something is missing in the definition of a hypothetical limiting frequency. (I don’t think this problem can be solved by changing the description of the experiment F . Someone might propose to add to the description F that “there is no robot sitting next to you that potentially intervenes with the tossing process,” yielding the description $F\&G$. But there may still be ways to manipulate $F\&G$ -tosses such that they always land heads. Given that a coin toss is deterministic, that is even likely.)

This paper won’t address the problem described here in detail, but I will simply assume that for any experiment F there is a *normal* way of repeating the experiment. Having a robot intervene when you are about to repeat a coin toss is an abnormal way of repeating coin tosses. If every normal way of repeating an experiment F leads to the same limiting frequency q , then we have $f_\infty^F(a) = q$; otherwise, the hypothetical limiting frequency does not exist.⁶

⁵ Lewis (1980) defines admissibility as follows: “Admissible propositions are the sort of information whose impact on credence about outcomes comes entirely by way of credence about the chances of those outcomes.”

⁶ Some might claim that different normal ways of repeating F lead to different limiting frequencies, in which case limiting frequencies as I defined them do not exist. Hájek (2009b), for example, claims that a fair coin can land heads forever. See Ackermans (forthcoming) for a proposed solution.

2.3 Applying a frequency principle

The frequency principles as formulated in this paper require the agent to know the limiting frequency. Since one typically doesn't know the limiting frequency for certain, the principles on their own are of limited use in practice. To apply a frequency principle in common scenarios of limited information, one needs to supplement it with additional assumptions or make it more general. Different schools of probability and statistics will disagree on the best way to do this. I briefly discuss how to do this in a subjective Bayesian fashion and an objective Bayesian fashion. The third example shows how one can use a frequency principle to test a theory of chance.

This section can safely be skipped.

Example 1: subjective Bayesianism

A subjective Bayesian would combine a frequency principle with the following principle of Bayesian updating.

Bayesian updating: Suppose E_1 and E_2 are mutually consistent sets of evidence. Then C is rational only if for any proposition a ,

$$C(a_1, \text{ given } E_1 \text{ and } E_2) = C(a_1 \mid E_2, \text{ given } E_1).$$

In words, one's credence in a_1 given the evidential situation described by E_1 and E_2 should be identical to one's conditional credence given the evidential situation described by E_1 , conditional on E_2 , as long as E_1 and E_2 are consistent. (Here conditional credence can be interpreted and defined in a variety of ways, as long as it satisfies $C(a \mid b)C(b) = C(a \text{ and } b)$, which allows us to manipulate it mathematically. See Sect. 2.2 for the used notational conventions.)

The Bayesian updating principle allows one to transform an evidential frequency principle as described above into a conditional frequency principle (which prescribes that one's credence, conditional on the limiting frequency being q , should be q). Such a frequency principle can be used even if one does not know the limiting frequency, on the basis of finite frequency data and a prior probability distribution of the value of the limiting frequency $C(f_\infty^F(a) \leq q)$. The procedure to calculate such a posterior probability is well-known and will not be discussed here (see Howson & Urbach, 2006, pp. 76–78).

Example 2: objective Bayesianism

This example uses an approach similar to the type of objective Bayesianism defended by Williamson (2010). Williamson's calibration norm (Sect. 3.3.1) is based on the principal principle, but its scope of application is increased, since it applies not just to evidence of precise chances, but also to evidence that the objective chance is within some interval or set (Sect. 3.3.1 in Williamson, 2010). Inspired by this approach, we can formulate a more general frequency principle as follows.

A Generalized Frequency Principle: Let $P \subset [0, 1]$, and let $\langle P \rangle$ be the closed convex hull of P . Suppose conditions Q hold. Then C is rational only if for all experiments F with outcome a and for all $i > 0$,

$$C(a_i, \text{ given } F_i \text{ and } f_\infty^F(a) \in P \text{ and } E) \in \langle P \rangle.$$

The gathering of the evidence $f_\infty^F(a) \in P$, as well as, possibly, verifying the conditions Q , will be left to frequentist statistics (cf. Williamson, 2010, pp. 43, 166–169). An objective Bayesian uses such evidence by combining the generalized frequency principle with the principle of maximum entropy to further restrict the set of rational credence functions. The following coin toss example illustrates this procedure.

Suppose that, after doing an experiment that is analyzed using frequentist statistics, one accepts the evidence that the objective chance of an F -toss resulting in heads lies in some interval $[0.4, 0.45]$. This evidence could imply that conditions Q are satisfied and that $f_\infty^F(\text{heads}) \in [0.4, 0.45]$ (see Sect. 5 for the relevant principles). By the generalized frequency principle, credence in heads should be in the interval $[0.4, 0.45]$. We now choose the credence function that maximizes entropy, which in this case is the function that assigns a credence to heads closest to 0.5. It follows that $C(\text{toss } i \text{ is heads, given } F_i \text{ and } f_\infty^F(\text{heads}) \in [0.4, 0.45]) = 0.45$.

Example 3: within a theory of objective probability

As many philosophers of chance believe, chances must play the role of guiding our credences. This role is typically described by the principal principle (see above). The principal principle can be used to test a philosophical theory of chance: if the theory says the chance is p , while it is irrational to set one's credence to p , then the theory must be false.

A frequency principle can be used to show that a particular theory of interest never fails this test. Let **FP** stand for whatever frequency principle one accepts, and let **T** stand for a particular metaphysical theory of chance that we aim to evaluate. Suppose our theory **T** assigns the chance $ch_F(a)$ to each outcome a of a chance setup F . Additionally, suppose that **T** satisfies the following Chance Frequency Condition, according to which chances must have the following relation to frequencies.

CFC: for all F, a , if $ch_F(a) = p$, then $f_\infty^F(a) = p$.

Now, suppose that the conditions Q of **FP** are satisfied whenever F is a chance setup and other available evidence is admissible (in the sense of the principal principle). Then, for any chance setup F , our frequency principle **FP** implies that it is rational to set your credences to the chances. Hence, the principal principle (for this particular theory of chance) follows from **FP**. In other words, if **FP** is correct, this theory of chance never violates the principal principle. This illustrates how a frequency principle can be used to establish that a particular theory of chance satisfies the key desideratum of being compatible with the principal principle.

3 The naive infinite frequency principle

The Naive Infinite Frequency Principle is the frequency principle with the least conditions. It states that if one knows only F , that each a_i is possible and not necessary (their credence is not 0 or 1), and that the hypothetical limiting frequency of a 's is q , then one's credence in a_i must be q .

NIFP: C is rational only if for all F , a and admissible evidence E such that $C(a_i, \text{ given } F_i) \in (0, 1)$, we have $C(a_i, \text{ given } F_i \text{ and } f_\infty^F(a) = q \text{ and } E) = q$.

(Here admissible evidence is defined differently from the principal principle, as evidence that rationality would allow you to disregard; see Sect. 2.2.) Some authors seem to defend **NIFP**, such as Mellor (1995, p. 45), although Mellor would perhaps claim to have defended a more restrictive version of the principle.⁷ In any case, it is instructive to point out what is wrong with it, in order to understand how the principle can be corrected.

The **NIFP** is naive because in certain situations it prescribes irrational credences. Hence, it doesn't satisfy the rationality desideratum. To see why this is the case, consider that limiting frequencies of sequences, by themselves, have no implications for finite initial segments of the same sequence, while the experimental description F might have such implications. For example, consider the sequence starting with 1000 zeros and ending with an infinite number of ones. The limiting frequency of ones in this sequence is 1, but the initial segment of the first 1000 entries has a frequency of 0. Or take the sequence $(s_n) = (1, 0, 1, 1, 0, 1, 1, 1, 0, \dots)$ with an increasing number of ones alternated with single zeros. This sequence has a limiting frequency of 1 of outcome '1', but any initial segment has a frequency of '1' lower than 1.

With this in mind, it is easy to come up with experiments that satisfy the conditions of **NIFP**, but for which it is irrational to set one's credence to the limiting frequency. The following example is about a pseudorandom number generator, a computer program that takes a seed (a number) as input and then deterministically outputs numbers—one number every time it is run—in such a way that the sequence it produces looks random. The limiting frequencies of the different outputs of pseudorandom number generators exist. Now let F be:

- *The program **RAND** is run once and its output recorded;*
- ***RAND** is a pseudorandom number generator that outputs either 0 or 1.*

Let a be the event that running F leads to output '1'. Given the information F , it could be reasonable to set

$$C(a_i, \text{ given } F_i \text{ and } f_\infty^F(a) = 1/2) = 1/2.$$

But now suppose that **RAND** is known to have a bias in initial segments, as implied by the experimental description G :

⁷ Mellor (1995, p. 45) asks the question: "Why should the evidence that some fact has a property which entails $f_\infty(P^*) = p$ justify the credence $cr(P) = p$?" Here P is a fact of type P^* . He then gives a betting argument that is supposed to defend this rule in the situation that $p \in (0, 1)$. This betting argument is only valid when it is assumed that credences in all facts of type P^* should be identical, but Mellor does not mention this condition.

- The program **RAND** is run once and its output recorded;
- **RAND** is a pseudorandom number generator that outputs either 0 or 1;
- The first 100 runs of **RAND** have a proportion of ones less than 0.5 (with certainty).

In this example, setting $C(a_i, \text{ given } G_i \text{ and } f_\infty^G(a) = 1/2) = 1/2$, as **NIFP** would require, is irrational for $i \leq 100$. With these credences, one would accept 100 bets whose payout is twice the price of the bet. That is, the actual value of each bet is a change in utility of $S(I(a_i) - 1/2)$, with $S > 0$. As before, S is the stake and $I(a_i)$ is the indicator function that returns 1 if a_i occurs and 0 if a_i does not occur. Note that the agent’s expected value of each bet, given her own credence function, is $E_C[S(I(a_i) - 1/2)] = 0$, and her expected value of 100 bets similarly is 0. However, as follows from the experimental description G , the amount of ones after 100 bets is less than 50. Hence, the actual value of these 100 bets together is

$$\sum_{i=1}^{100} S(I(a_i) - .5) = S(\sum_{i=1}^{100} I(a_i) - 50) < 0,$$

a sure loss. The **NIFP** does not satisfy the rationality desideratum.

One could object that G should be seen as inadmissible evidence. However, since **NIFP** does not contain an admissibility requirement for the experimental description (only for additional evidence E), **NIFP** still entails a credence of 1/2 in the above example, and so it entails accepting the 100 bets for which there is a sure loss, as I argued above. So this objection, as it stands, does not refute my argument against **NIFP**. However, the objection does get something right. There is a closely related principle, with a more extensive admissibility requirement, for which the counterexample does not hold. This objection will be discussed in detail in Sect. 3.1 below.

Another objection to the above example could be that the *normal* way of repeating the experiment G is by resetting the computer’s memory such that the pseudo-random number generator will output its first number each time. (Note that normality refers not to the description G but to the way in which the conditions G are recreated each time—see the end of Sect. 2.2.) Such a conception of normality is different from the one that I have been using so far in this paper. This alternative conception leads to a different frequency principle, which applies if one has knowledge of a limiting frequency as defined with respect to this alternative conception of normality, as opposed to limiting frequencies defined with respect to our original conception of normality. Hence, such an objection can best be seen as proposing a different version of **NIFP** that is claimed to be more plausible.

In any case, this objection leads to other problems. Since it is not clear how the program’s seed is generated, it is not clear what the limiting frequency of this procedure is and whether it exists at all. Suppose that the seed is read from a persistent memory. In that case, G will output the same number each time, so the limiting frequency cannot be 1/2 (it can only be 1 or 0). If it is generated by some deterministic process, but a different one each time, then *that* process could still lead to different frequencies in the short run than in the long run (we could repeat the above argument). Only if the seed is indeterministically generated—say, using a quantum computer—it is clear that the frequencies of outcomes will be similar in the short run as in the limit.

However, to satisfy the informativeness desideratum, the frequency principle should have something to say about deterministic cases. Hence, it seems that this objection ultimately leads to the conclusion that **NIFP** does not satisfy either the rationality desideratum or the informativeness desideratum.

For those skeptical about the above example, here is another example with coin tosses. Suppose we repeat tosses of a biased coin, which comes up heads about a third of the time initially, but has limiting frequency $1/4$. The experiment is described by F :

- *A biased coin with characteristics C is tossed;*
- *A C -coin's frequency of landing heads after tossing it a large amount, but feasible within a human lifetime, is likely very close to $1/3$.*

With this experiment, supposing a is the event that an F -toss lands heads, **NIFP** implies $C(a_i)$, given F_i and $f_{\infty}^F(a) = 1/4 = 1/4$. However, it is intuitively clear that to avoid losses in betting games, one should set one's credence in 'heads' to $1/3$ rather than $1/4$. In this example there are no sure losses—they are only "likely"—so application of the sure loss principle is not possible. Nevertheless, similar intuitions point towards rejecting **NIFP** in this example.

Like above, someone might object that it should not be possible for an experiment described by F to have a limiting frequency of $1/4$ if it is repeated *normally*—in a different sense of 'normal' than I use it. What could such a different conception of normality be? If it is to repeat the exact same conditions as the first toss each time—that is, repeat not just F but also all the other conditions surrounding the experiment when it was first run—then the limiting frequency must be 0 or 1. A frequency principle based on such a conception of normality would only be useful for indeterministic experiments. According to yet another version of normality, one imagines a way to repeat the coin toss such that the limiting frequency is actually $1/3$ rather than $1/4$ as above. The only way in which the limiting frequency could be $1/4$ (one could reason) is if outside conditions are gradually changing over time in a way such that the bias of F -tosses asymptotically changes from $1/3$ towards $1/4$. The solution would be to think of *normal* repetitions as fixing certain conditions of the world—those conditions that cause the C -coin to land heads about a third of the time initially—but not others. Call this conception of normality *stable-normality*. When using stable-normality, the quantity $f_{\infty}^F(a)$ would refer to the limiting frequency if a C -coin were repeatedly tossed, but other conditions of the world would be fixed (or prevented from changing) in such a way that its bias remains the same.

While this solution looks satisfactory in this particular example, a frequency principle based on stable-normality could lead to irrational choices in general. A frequency principle advises an agent who bets on F -experiments that will actually be repeated, but actually repeating an F -experiment is very different from stable-normally repeating an F -experiment. Take the above example but now suppose that a C -coin's bias changes more quickly, such that after 1000 tosses the frequency of 'heads' is already likely to be closer to $1/4$, as we suppose follows from an experimental description " G ". We suppose the agent's evidence consists of G and $f_{\infty}^G(a) = 1/3$ (defined with respect to stable-normality). It seems irrational to have a credence of $1/3$ in 'heads', since the agent would take these 1000 bets which have a frequency of 'heads' close to

1/4. But since the stable-normal limiting frequency is 1/3, **NIFP** would recommend a credence of 1/3. In this situation, the **NIFP** based on stable-normality looks irrational.

3.1 Intermission: an admissibility paradigm

As I suggested above, one could object to the above examples that the experimental descriptions are themselves inadmissible evidence. Such an objection is best understood as favoring a slightly different version of **NIFP** which I will call the Simple Infinite Frequency Principle (**SIFP**). Roughly speaking, the **SIFP** moves as much information as possible out of the experimental description F and the frequency evidence $f_{\infty}^F(a)$ and into the admissible evidence portion of the credence function, E . The upshot of this objection is that one could also frame the problems discussed in this paper in terms of the question of what evidence is admissible with respect to frequency principles, although I do not choose to do so.

Imagine a situation in which an agent solely has the following information: (a) there is an event type x such that the outcome event a_i is preceded by an experiment event x_i of type x and x_i occurs; (b) the limiting frequency of outcome type a when experiments of type x are repeated indefinitely is $f_{\infty}^x(a) = q$; (c) any other admissible evidence E . For example, a_i could be the event that the experiment denoted x_i results in outcome type ‘heads’ (a). E might consist of the information that x is a normal toss of an unbiased coin. (In the context of a frequency principle, one must have the information that x_1 occurs—that is, the experiment that would precede a_1 is actually performed. If the agent does not know whether x_1 will be performed, her credence in a_1 will depend on her credence that x_1 will be performed, since a_1 will not occur if x_1 is not performed; in that case, her credence in x_1 cannot depend on the limiting frequency alone, as an infinite frequency principle dictates.)

Note that an agent with the credence $C(a_1, \text{given } x_1)$ does not know anything about the experiment type x . All that the agent with this credence knows is: there is an experiment type denoted x ; such an experiment x can result in outcome type a ; and an experiment of outcome type x , denoted x_1 , is actually performed.

We are now in a position to state **SIFP**.

SIFP: C is rational only if for all experiment types x with outcome type a and admissible evidence E , we have $C(a_i, \text{given } x_i, f_{\infty}^x(a) = q \text{ and } E) = q$.

In the case one knows something about the experiment type x , such as x is a coin toss with an unbiased coin, we could denote that information as Fx (read: x satisfies the property F). Whether **SIFP** requires us to set $C(a_1, \text{given } x_1, f_{\infty}^x(a) = 0.5 \text{ and } Fx) = 0.5$ now depends on whether Fx is admissible evidence.

SIFP is able to evade susceptibility to the counterexamples to **NIFP**, making it more plausible. If the admissibility requirement for E is sufficiently strong, then it is hard to see how counterexamples against **SIFP** could exist. The situation in which one does not have any additional evidence E is too sparse to allow for the creation of counterexamples, which typically involve additional evidence. When one does have additional evidence E , it can always be claimed that E is inadmissible.

However, as long as ‘admissibility’ is not clearly defined, **SIFP** is not very informative (in comparison to the other principles discussed in this paper). The definition I gave above, adapted to the current paradigm, is as follows: admissible evidence is evidence that rationality would allow you to disregard when x_i and $f_\infty^x(a) = q$ are also part of your evidence. The **SIFP** would be informative in the minimal but uninteresting case in which the agent has no additional evidence. In any interesting and realistic situation, the agent will have additional evidence E . To apply the principle in such cases, we need to be able to ascertain whether E is admissible. But note that admissibility is defined in terms of rationality, and that the purpose of a frequency principle is precisely to find practicable conditions for rationality. Hence, as long as admissibility is not given further analysis, the **SIFP** is uninformative in all interesting cases. In order for **SIFP** to be useful, one still needs to figure out, at the very least, which experimental descriptions Fx are admissible. That task is the same as the task undertaken in this paper, although it is now framed in terms of admissibility. The **CIFP** and **EIFP** introduced in Sects. 5 and 6, respectively, are more informative frequency principles due to their explicit conditions for ‘admissible’ experimental descriptions F .⁸

4 Positive arguments for frequency principles

Before turning to the other principles, I discuss two arguments that can be used to defend infinite frequency principles.

Most arguments for infinite frequency principles offered in the literature seem to be betting arguments (Mellor, 1995; Howson & Urbach, 1993; Williamson, 2010), but they suffer from serious defects. Such arguments rely on an often overlooked assumption that all repetitions of the experiment must be given equal credences, which needs a separate defense (see also Strevens, 1999; Albert, 2005). However, as I show below, even under the equal credence assumption the betting arguments offered in the literature do not seem to be valid. The problem is that convergence to a sure loss in the limit to infinity is no ground for rejecting a credence function. The popularity of these arguments is peculiar, since the equal credence assumption allows for a simpler consistency argument, that I give in Sect. 4.2.

As I show, the betting argument doesn’t prove in a mathematically precise way what it claims to prove. Nevertheless, it does lead to some weaker results that might lend some plausibility to the thesis that, given that equal credences are required, an agent who does not calibrate to the limiting frequency is likely to face losses. A similar conclusion can be drawn from the consistency argument. While the consistency argument establishes no conclusion about betting losses, it does establish a violation of the probability axioms. If the probability axioms are necessary conditions of rationality, as is plausible, it follows that setting one’s credence to the limiting frequency in a situation of equal credences is a necessary condition of rationality.

⁸ Note that this problem does not apply to the same extent to the principal principle. The principal principle applies to interesting situations in which one’s only evidence is an outcome’s chance and the fact that this outcome follows a chance experiment. Knowledge that one faces not just an experiment, but a *chance* experiment, is significantly more specific than the minimal situation in which **SIFP** applies.

An important feature of both arguments is that they restrict attention to credence functions which assign equal credences to all repetitions of the experiment F . That is, C satisfies, for some $q \in [0, 1]$ and all $i > 0$,

$$C(a_i, \text{ given } F_i \text{ and } E) = q. \tag{1}$$

The subsequent sections will consider the conditions under which an agent is rationally required to assign equal credences to all repetitions of an experiment.

4.1 The betting argument

The betting argument, as it is commonly formulated, is as follows. Let F describe a repeatable experiment and let a be an outcome of F . Suppose that the evidence for each run of the experiment is restricted to F_i , that $f_\infty^F(a) = p$ and that the equal credence condition is satisfied. We also suppose that the agent is in a situation in which it is rational for her to make any number of bets on F -experiments.

Now suppose that the credences are unequal to the limiting frequency, that is, for all $i > 0$ we have $C(a_i, \text{ given } F_i \text{ and } f_\infty^F(a) = p) = q \neq p$. Let $f_n^F(a)$ be the frequency of a 's after n repetitions of F . By the definition of a limit, there is some $N \in \mathbb{N}$ such that for all $n > N$, we have $|f_n^F(a) - p| < |q - p|$. We consider (without loss of generality) only the case of $p < q$ and the betting stake $S = 1$. Our agent would accept (in suitable situations) the n bets whose actual value is $I(a_i) - q$, for $i = 0 \dots n$. Let $n > N$. Then the total value of these n bets is

$$\sum_{i=1}^n I(a_i) - q = n(f_n^F(a) - q) < n(p + |q - p| - q) = 0. \tag{2}$$

Hence, the agent whose betting quotients are as assumed will have losses if she bets N or more times. (If $p > q$ a similar argument can be made.)

Most authors conclude the argument here, as if it is now clear that C is irrational.⁹ However, the argument does not seem to establish anything significant. There is no sure loss for the agent for any number of bets, since the *value* of N is not implied by her evidence. For every number $n \in \mathbb{N}$, it is not a sure thing that $n > N$, and therefore, there is no sure loss after n bets. Hence, it seems to me that this way of putting the argument does not lead to the desired conclusion. (Of course, the agent's evidence implies that there *exists* some unknown number N after which she faces a loss, but this does not mean that it is irrational to bet N times, as the following example shows. If a fair coin is tossed indefinitely, then there exists a number of tosses M for which the frequency of 'heads' is greater than $1/2$ —in fact, there are infinitely many such numbers. But it is not irrational to bet on a fair coin M times with betting quotient $1/2$.)

⁹ Williamson (2010, p. 41), for example, shows that there exists some N such that for $n > N$, the agent faces a sure loss, and concludes, "Hence, the agent is sure to lose money in the long run." If "in the long run" means "after a countably infinite number of bets", the statement is uninteresting, since no one can ever take an infinite number of bets. On the other hand, if "in the long run" means "after a very large number of bets", the statement is false: there is no sure loss after any number of bets, however large.

One might alternatively try to argue that the agent is very close to being sure of a loss given her own credence function. The first step in this argument is to note that $C(f_\infty^F(a) = p, \text{ given } f_\infty^F(a_i) = p) = 1$, that is, a_i converges to p almost surely. It follows from this that $f_n(a)$ converges in probability to p .¹⁰ Let P be the probability function given by $P(\phi) = C(\phi, \text{ given } f_\infty^F(a_i) = p)$. For all $\varepsilon > 0$ we have

$$\lim_{n \rightarrow \infty} P(|f_n(a) - p| \leq \varepsilon) = 1. \tag{3}$$

It is then easily shown that the agent’s credence that she faces a loss converges to 1, that is,

$$\lim_{n \rightarrow \infty} P\left(\sum_{i=1}^n I(a_i) - q < 0\right) = 1. \tag{4}$$

This means that the agent will become almost completely sure that she faces a loss when n becomes large. However, almost sure is not sure, and accepting bets that almost certainly lead to sure losses is not necessarily irrational. To see this, consider the following example: one bets on an outcome a that is very unlikely, say $C(a) = 0.001$. Suppose the price of this bet is \$10 and the payout if a occurs is \$100,000. The expected value of this bet is \$90, but one is almost sure of a loss of \$10. Such a bet can be rational to take.

Hence, this approach also leads us nowhere. The versions of the betting argument discussed here do not lead to the desired conclusion that a credence function that does not satisfy a frequency principle is irrational.

One could proceed by showing that the agent also assigns a negative expected value to a sufficiently large number of bets, and therefore would both accept and not accept these bets. However, such an argument no longer has the structure of a betting argument but looks more like a consistency argument.

4.2 The consistency argument

Let P be the probability function given by $P(\phi) = C(\phi, \text{ given } f_\infty^F(a_i) = p \text{ and } F_i, \text{ for all } i > 0)$. Again, we assume equal credences, that is, $P(a_i) = q$. The argument starts in the same way as the second betting argument above. As is commonly accepted, we have $P(f_\infty^F(a_i) = p) = 1$. It follows from this that $f_n(a)$ converges in probability to p ,¹¹ that is, for all $\varepsilon > 0$ we have

$$\lim_{n \rightarrow \infty} P(|f_n(a) - p| \leq \varepsilon) = 1. \tag{5}$$

Suppose that $p \neq q = P(a_i)$. We only treat the situation that $p > q$. Let $\varepsilon > 0$ such that $(p - \varepsilon)(1 - \varepsilon) > q$ and $\varepsilon < p$. Let N_ε such that $P(|f_{N_\varepsilon}(a) - p| \leq \varepsilon) > 1 - \varepsilon$.

¹⁰ For a proof, see Theorem B.90 in Schervish (1995, p. 638).

¹¹ See footnote 10.

Taking the mathematical expectation with respect to P of the frequency $f_{N_\varepsilon}(a)$,¹² we get

$$\mathbf{E}[f_{N_\varepsilon}(a)] \geq \mathbf{E}[f_{N_\varepsilon}(a) \mid (|f_{N_\varepsilon}(a) - p| \leq \varepsilon)] \cdot P(|f_{N_\varepsilon}(a) - p| \leq \varepsilon). \quad (6)$$

The first term on the right hand side of (6) is at least $(p - \varepsilon)$, and the second term is greater than $(1 - \varepsilon)$, so we have $\mathbf{E}[f_{N_\varepsilon}(a)] > q$. At the same time, by elementary probability calculus, we have $\mathbf{E}[f_{N_\varepsilon}(a)] = \sum_{i=1}^n \mathbf{E}[I(a_i)]/n = q$, in contradiction. So $p \leq q$, and by a similar argument, $p \geq q$.

This argument shows that for any credence function C that satisfies the probability axioms and assigns equal credences to all a_i given $f_\infty^F(a) = p$, we have $C(a_i, \text{ given } f_\infty^F(a) = p) = p$. Note that the latter statement, understood as a frequency principle, does not satisfy the informativeness desideratum, since it does not constrain the set of rational credence functions further than do the probability axioms (all equal-credence functions that satisfy the probability axioms satisfy this frequency principle). What remains to be found are external conditions under which rationality *requires* equal credences. This question will be explored in the next two sections.

It is worth pointing out that another and well known consistency argument exists given the more stringent condition that the a_i are exchangeable. The set of Bernoulli variables $X_1 = I(a_1), X_2 = I(a_2), \dots$ is exchangeable when for each finite subset X_{i_1}, \dots, X_{i_n} , values $j_1, \dots, j_n \in \{0, 1\}$ and permutation σ of the indices the we have $P(X_{i_1} = j_1, \dots, X_{i_n} = j_n) = P(X_{\sigma(i_1)} = j_1, \dots, X_{\sigma(i_n)} = j_n)$.¹³ Under the condition of exchangeability and a probability function P that satisfies the probability axioms, De Finetti’s representation theorem implies that we have $P(a_i \mid f_\infty^F(a) = p) = p$ (as well as that the a_i are conditionally independent, conditional on the limiting frequency).¹⁴

To defend a frequency principle on the basis of the representation theorem, we would have to state conditions Q under which exchangeability is rationally required. Exchangeability, however, is a demanding condition. Equal credences in the a_i is a less demanding condition, which will both make a frequency principle more general and easier to defend.

5 The chancy infinite frequency principle

The Chancy Infinite Frequency Principle has as its condition that F should be chancy. It states that if one knows that F is chancy and that the limiting frequency of a is q , then one’s credence in a must be q . I will say that an experiment is chancy if and only if the objective chance of each outcome a given F , written $ch_F(a)$, exists.

CIFP: C is rational only if for all F , a and admissible evidence E we have $C(a_i, \text{ given } F_i, F \text{ is chancy, } f_\infty^F(a) = q \text{ and } E) = q$.

¹² For a discrete random variable X with possible values x_1, \dots, x_k , the mathematical expectation is defined as $\mathbf{E}[X] = \sum_{i=1}^k x_i P(X = x_i)$.

¹³ $P(a, b, c)$ is the probability of the conjunction $P(a \wedge b \wedge c)$.

¹⁴ See Schervish (1995) for a modern presentation of the representation theorem.

Since there are multiple theories of chance, **CIFP** is better thought of as a collection of principles, one for each definition of ‘chanciness’. Not all of these principles may be true or satisfy our desiderata.

When chances are interpreted as limiting frequencies in the sense of Mises (1957), you get the version of **CIFP** endorsed by Howson and Urbach (2006). Strevens (1999) and Williamson (2010, pp. 39–42) implicitly discuss a version of **CIFP** without committing to a particular interpretation of chance, but instead adopt the Chance Frequency Condition (called the *long run frequency postulate* by Strevens (1999)).¹⁵ The Chance Frequency Condition (**CFC**) is a plausible condition on chances that states that type chances should be identical to limiting frequencies.¹⁶

CFC: for all F, a , if $ch_F(a) = p$, then $f_\infty^F(a) = p$.

If we assume **CFC**, then **CIFP** is equivalent to the principal principle:

PP: C is rational only if for all F, a and admissible evidence E we have $C(a_i, \text{given } F_i, ch_F(a) = p \text{ and } E) = p$.

Hence, if one accepts both **CFC** and the principal principle, then one should accept **CIFP**.

The argument put thus is valid, but insufficient as grounds to accept **CIFP** (as a principle that is both true and satisfies our desiderata). First, **PP** itself is often justified by an appeal to a principle like **CIFP** (see Strevens, 1999). In that case, the equivalence of **CIFP** and **PP** is no justificatory ground for accepting **CIFP**. Hence, we need a separate argument to justify **CIFP**, such as one based on the arguments in the previous section.

Second, whether **PP** is true, and whether **CIFP** satisfies the rationality desideratum, depends on one’s theory of chance. (I remind the reader that the rationality desideratum requires that the principle constrains in the right way, i.e., that it scraps only irrational functions.) To interpret **PP**, one needs a theory of chance, and different theories of chance will lead to different versions of **PP**, not all of them which will be true. One way in which a version of **PP** could be false is if it depends on a conception of chanciness (of an experiment F) that does not imply equal credences for a_i given F_i . Recall that the examples of Sect. 3 exploit the irrationality of having equal credences in certain situations to derive the irrationality of setting one’s credence to the limiting frequency. If chanciness does not imply equal credences, it is conceivable that similar examples of irrationality exist in which the experiment is chancy. In such cases, using **CIFP** to calibrate one’s credences would be irrational. On the other hand, if the chanciness of F does imply the equal credence condition, then the arguments in the previous section

¹⁵ Williamson defends **CIFP** only implicitly, while arguing for his version of the principal principle that uses single case chances, called C2. This argument uses the betting argument to defend **CIFP**, but establishes C2 via an implicit argument that connects **CIFP** with C2. Similarly, while Strevens discusses critically a betting argument that would defend **CIFP**, it is intended to defend a single-case version of the principal principle.

¹⁶ Note that **CFC** is stronger than the strong law of large numbers. It is possible to reject **CFC** without getting into conflict with the probability axioms. While I’m not aware of authors who explicitly argue against it, some of Hájek’s (2009b) arguments against hypothetical frequentism can be seen as arguments against **CFC**.

make it plausible that **CIFP** satisfies the rationality desideratum. The question is, then, which conceptions of chanciness imply the equal credence condition.

Third, whether the informativeness desideratum is satisfied also depends on what chanciness is, as **CIFP** is informative only if deterministic chances can be non-trivial (different from 0 or 1). If deterministic experiments cannot have non-trivial chances, then the chance of a coin toss landing heads can only be 0 or 1, which leads to the following situation. Suppose that F describes a (deterministic) coin toss. If F is chancy, then, by **CFC**, the limiting frequency must be either 0 or 1. There might be a sense of ‘limiting frequency’ for which the latter statement is true, but the value (0 or 1) of a limiting frequency in this sense is not typically part of our evidence. Hence, **CIFP** would be uninformative. (Conversely, if the limiting frequency is $1/2$, then by **CFC**, F is not chancy; hence, **CIFP** would have nothing to say in this situation.)

Hence, we look for a conception of chanciness that implies the equal credence condition and ensures that the informativeness desideratum is satisfied.

A conception of chanciness does not need to be based on a precise definition. Chanciness could be left undefined and understood as a metaphysical primitive. One could then just assume that chances satisfy **PP**, **CFC** and the equal credence condition. However, there are substantive disagreements about which experiments are chancy and not—for example, there is the disagreement about whether deterministic experiments can be chancy—to such an extent that it is unclear whether there is a single underlying concept of chanciness. Hence, the ‘metaphysical primitive’ interpretation of **CIFP** does not satisfy the precision desideratum. (Recall that the precision desideratum requires that a principle is not ambiguous, i.e., open to various different interpretations.)

Similarly, one could use a no-theory theory of chance as proposed by Sober (2010). The no-theory theory argues that chance, understood as a primitive, can be assumed to be an objective quantity because different measurement procedures lead to the same results. If this is indeed the case—if there is agreement on what the chances are *and* these include non-trivial deterministic chances, as Sober intends—then the above problem of deterministic cases does not arise. However, a different sort of problem may remain. Some authors have expressed skepticism that the principal principle can be justified under a no-theory theory (Hájek, 2007; Hoefer, 2007). If all that we know about chances is how to measure them (and that these measurement procedures converge), then why should we believe that these chances are the right sort of thing to guide our credences? That is to say, why should we believe that these chances satisfy the principal principle and the equal credence condition? The no-theory theory may work well as an interpretation of probabilistic explanation as used in evolutionary theory and statistical mechanics; but without further justification, we cannot be sure that no-theory chances warrant the use of chances to guide our credences.

An alternative approach is to give a reductive definition of ‘chanciness’. While a version of **CIFP** based on a precise definition satisfies the precision desideratum, it is unclear whether the other two desiderata can be satisfied. I discuss this approach below.

5.1 Defining chanciness

We need a definition of chanciness that (1) can be shown to imply the equal credence condition and (2) allows for the possibility that deterministic chance experiments have non-trivial chances. A version of **CIFP** based on such a concept of chanciness plausibly satisfies the rationality desideratum and informativeness desideratum. I argue that two prominent candidates—maximal specificity and von Mises/Church randomness—do not have the desired properties. It does not seem to be the case that precise necessary and sufficient conditions for chanciness have been offered in the literature that do have these properties. This means that anyone who endorses **CIFP** should endorse an uninterpreted conception of chance, at least until a better definition of chanciness is found. While this could lead to a failure to satisfy the precision desideratum, that may be considered more acceptable than a failure of the other two desiderata.

Note that the arguments in this section might also lead one to conclude that a chance coordination principle using type chances is undesirable, and one should instead adopt a principle of direct inference that uses single case chances, such as the original principal principle as proposed by Lewis.¹⁷ Such a principle is of less relevance to infinite frequency principles, so this line of argument will not be discussed here.

One potential definition of chanciness could involve the concept of maximal specificity as used by propensity interpretations of probability. A description F of an experiment is maximally specific when there is no description G containing additional characteristics of the experiment, such that $F \& G$ correctly describes this experiment and the limiting frequency of $F \& G$ -experiments is different from the limiting frequency of F -experiments. For example, suppose F describes a coin toss whose behavior is as described by Suppes's model, and suppose that G describes, for a given coin toss, all of the toss's characteristics, such as upward velocity and spin. (Recall that in Suppes's model, the way a coin lands is fully determined by the coin's characteristics like upward velocity and spin.) If the limiting frequency of F -tosses is $f_{\infty}^F(\text{heads}) = 1/2$, then F is not maximally specific, since the limiting frequency of $F \& G$ -tosses, $f_{\infty}^{F \& G}(\text{heads})$, is either 0 or 1. On the other hand, $F \& G$ is maximally specific, since no details could be added that would change the limiting frequency.

We could use maximal specificity as a defining condition for an experiment to be *chancy*. But the downside of a definition of chanciness based on maximal specificity is that it rules out that deterministic chance experiments have non-trivial probabilities (see also Fetzer, 1981).¹⁸ A frequency principle based on such a definition of chanciness would have nothing to say about deterministic experiments with non-trivial limiting frequencies, such as coin tosses. Hence, it would not satisfy the informativeness desideratum.

Another potential definition of chanciness could involve the definition of randomness as used within hypothetical frequentism, an approach similar to Howson and Urbach (2006). A sequence of outcomes is said to be *von Mises/Church ran-*

¹⁷ Thanks to a reviewer for pointing this out.

¹⁸ Of interest are long-run propensity views that do not accept the principle of maximal specificity. It has been argued that these allow for non-trivial deterministic chances (Niiniluoto, 1981; Mellor, 2005; Berkovitz, 2015). However, as far as I'm aware they do not offer a precise alternative that could play the role of defining 'chanciness'.

dom if and only if there is no computable algorithm that selects a subsequence with a different limiting frequency than the original sequence, taking only previous outcomes as input (Mises, 1957; Church, 1940). A more formal definition of randomness is offered by Church, but here I explain the concept using an example. Suppose a repeated experiment outputs a sequence with a predictable pattern, like $(x_n) = (1, 0, 0, 1, 0, 0, 1, 0, 0, \dots)$. This sequence has limiting frequency $1/3$ of outcome '1'. There is an algorithm that selects the subsequence $(x'_n) = (1, 1, 1, \dots)$ by taking every third member of (x_n) . Since (x'_n) has the different limiting frequency 1 of outcome '1', (x_n) is not random.

An experiment F could be defined to be chancy if and only if F produces random sequences when repeated indefinitely. Unfortunately, this definition of chanciness does not satisfy the rationality desideratum, as random sequences may behave very differently in the limit than they do in finite initial segments (in which case the equal credence condition is not satisfied). For example, one could create a random sequence using a random number generator that outputs '0' or '1', such that for the n 'th number in the sequence the probability of '0' is chosen to be $1/(n+1)$. It seems plausible that random number generators exist such that this sequence is von Mises/Church random (for example, those that get their randomness from a quantum computer). It has limiting frequency 1 of output '1', but its initial segments are likely to contain many zeros. Suppose F describes such a random number generator. That is, F is given by:

- The program **RAND2**, a pseudorandom number generator, is run;
- The objective probability of output '0' when **RAND2** is run for the n 'th time is $1/(n+1)$;
- Sequences produced by repeatedly running **RAND2** indefinitely are von Mises/Church random.

Given a suitable notion of objective probability, the limiting frequency of ones is $f_\infty^F(1) = 1$, so **CIFP** would require $C(\text{output '1'}, \text{given } F_i) = 1$. However, for any finite amount of repetitions of F it is possible that a '0' is outputted once or more. In this case, setting one's credence of output '1' to 1—hence accepting bets with betting quotient $q = 1$ —does not lead to a sure loss, as it is still possible that F outputs only ones. However, this is a case in which one surely does not gain and a loss is possible, which seems irrational all the same. If F outputs only ones, then taking bets with betting quotient $q = 1$ has a payoff of 0; if F outputs at least one zero, then the payoff is less than 0. Another reason that the chosen credence must be irrational is that the agent should deem it possible that a zero is outputted, and should therefore assign a non-zero credence to that event. Hence, under an interpretation of **CIFP** based on von Mises/Church randomness, it does not satisfy the rationality desideratum.

Various attempts have been made to create stronger definitions of randomness that do not suffer from the problems of von Mises/Church randomness (see Eagle, 2016). One example is Martin-Löf randomness. Martin-Löf randomness implies von Mises/Church randomness, but it has implications for finite initial segments as well. The latter feature is both its strength as well as its weakness. Martin-Löf randomness may not suffer from some of the counterexamples regarding initial segments that apply to von Mises/Church randomness. However, actual chancy experiments occasionally

produce non-random outcome sequences. It is possible, for example, that a coin lands heads a thousand times in a row. A ‘chanciness’ defined using Martin-Löf randomness would exclude almost all experiments—including coin tosses—that are commonly considered chancy, and therefore would not satisfy the informativeness desideratum.

Given that existing definitions of chanciness don’t work, one might instead leave chanciness uninterpreted, as discussed above, and rely on an intuitive or no-theory understanding of chanciness. In the examples in Sect. 3, there is always a sense in which the experiment is intuitively non-chancy. A chancy experiment’s outcomes have the same chance every time they are run, but these computer programs’ outcomes do not, one might intuitively think. Hence, these are not obviously counterexamples to the uninterpreted **CIFP**.

Since the uninterpreted **CIFP** does not satisfy the precision desideratum, it might be considered undesirable. For those who, like myself, value precision, **CIFP** can be saved only if new research shows that chanciness is reductively definable after all. However, even if that project fails, not all hope is lost, as I show in the next section that a different sort of frequency principle exists that satisfies all desiderata.

6 The equal credence infinite frequency principle

As the consistency argument in Sect. 4.2 shows, it follows from just the probability axioms that when credence functions are restricted to functions that assign equal credences to repetitions of the experiment, these functions all set their credence to the limiting frequency. Hence, if the probability axioms are necessary requirements of rationality (as is plausible) and one’s frequency principle restricts itself to situations in which equal credences are a requirement of rationality, then that frequency principle satisfies the rationality desideratum (i.e., does not constrain too much). If we find conditions under which equal credences are rationally required, we can formulate a frequency principle that applies in just the right cases. I introduce such a principle, which I argue satisfies all desiderata. It is the only principle that satisfies the informativeness and rationality desideratum as well as the precision desideratum, as it does not rely on a metaphysical primitive. Moreover, it satisfies the informativeness desideratum to a greater extent than the other frequency principles discussed so far.

A common argument for identical credences uses what we may call the equal evidence postulate, which states: if all one knows about each a_i is that it is an outcome of an experiment F with limiting frequency $f_\infty^F(a) = q$, then one ought to have identical credences in the a_i .¹⁹ (See Howson & Urbach, 1993, p. 345; Williamson, 2010, p. 41). But note that these authors only argue for equal credences as part of their defense of a version of **CIFP/PP**. Hence, they consider the equal evidence postulate when applied to repetitions of *chance* setups. When restricted to chance setups, the equal evidence postulate is more plausible. However, the arguments they offer do not depend on this restriction, so I discuss these arguments in isolation.) There is some room for interpretation of this postulate. On a liberal interpretation, it is false. On a

¹⁹ This tradition of symmetry principles used to derive a frequency principle extends to logical probability. See Carnap (1962, Chapter VIII).

strict interpretation, its application is limited and it cannot be used to defend an infinite frequency principle.²⁰ Although I don't accuse anyone of this, one can sometimes get the impression that authors defend the plausible strict version, while the useful liberal version is applied in an argument for a frequency principle.

On the liberal interpretation, an agent is considered to have equal evidence for the a_i even if evidence about the a_i that is normally taken to be implied by the description of a_i is different—call this *weakly equal evidence*. For example, when speaking of an agent betting on the event described as “the i 'th toss of an unbiased coin lands heads”, it is normally assumed (as I do elsewhere in the paper) that it is part of the agent's evidence that i unbiased coins are tossed preceding that event. An agent whose evidence about a_i includes that an F -experiment has been conducted $i - 1$ times before one gets to the F -experiment that precedes a_i , but whose other evidence for the different a_i is identical, is still considered to have weakly equal evidence. The liberal equal evidence postulate states that an agent with weakly equal evidence in the a_i should assign equal credences. However, as demonstrated by examples elsewhere in this paper (Sects. 3 and 5.1), evidence that F -experiments have been conducted a number of times before can require a rational agent to assign different credences to different a_i . Hence, this version of the postulate is false and thus cannot assist in defending an infinite frequency principle.

On the second, strict interpretation, an agent is considered to have unequal evidence if it includes that an F -experiment has been conducted $i - 1$ times before one gets to the F -experiment that precedes a_i —call this *strict equal evidence*. The strict equal evidence postulate states that an agent whose evidence for the a_i is strictly equal should assign equal credences. The strict equal evidence postulate looks plausible to me, and Williamson (2010, p. 41) gives some considerations against the intuitive tenability of violating a strict equal evidence postulate. However, the strict equal evidence postulate is not able to support an infinite frequency principle. To see why, note that under strictly equal evidence, the agent does not know that a_2 is a later event than a_1 —otherwise, the agent would also have to know that an F -experiment is conducted twice before a_2 , and the evidence would not be strictly equal. This implies that, from the perspective of the agent, the sequence of events a_1, a_2, \dots might be ordered differently than the time-ordered sequence of the actual experiments. The limiting frequency $f_\infty^F(a)$, however, is based on the time-ordered sequence of outcomes, and a differently ordered sequence might have a different limiting frequency (see also Hájek, 2009b). As a result, when one takes the (possibly) non-time-ordered sequence a_1, a_2, \dots and the time-ordered limiting frequency $f_\infty^F(a)$ as input of the betting argument or the consistency argument, neither argument works. In case of the betting argument, betting on the sequence a_1, a_2, \dots (where the bets are settled in that order rather than the time-order of events) no longer implies a sure loss, since the betting argument derives a sure loss from the time-ordered limiting frequency.²¹ In case of the consistency argument,

²⁰ See also Strevens (1999, pp. 261–262), who provides arguments against an equal evidence postulate that may apply to both its strict and liberal version.

²¹ One might object that if the agent is willing to bet on any individual a_i , she would be willing to bet on any sequence of a -events regardless of order. Hence, she would also bet on the sequence $a_{\phi(1)}, a_{\phi(2)}, \dots$, where $\phi(i)$ permutes the agent's sequence to the time-ordered sequence. However, if the agent does not know that the permuted sequence of bets under consideration is the time-ordered sequence, she does not

Eq. (5) in Sect. 4.2 does not hold due to the different orderings of the agent's and the time-ordered sequence. (Williamson, 2010, p. 41, offers a version of the strict equal evidence postulate that, on one interpretation, might be thought to be superior to the version discussed here, but it suffers from a similar problem as the liberal equal evidence postulate.²²)

Nevertheless, the equal evidence postulate has some intuitive appeal, and it stands to reason that a similar principle will do the job. I suggest that equal credences are a requirement of rationality, except when unequal credences are a requirement of rationality. When counterexamples to equal credences (such as those offered in Sect. 3) exist, rationality requires unequal credences. When no such counterexamples exist and existing principles of rationality leave open both equal and unequal credences, we let in the intuitive appeal of the equal evidence postulate and say that equal credences are required. This is the Equal Credence Principle.

ECP: Let F be some experiment with outcome a for which it is “*prima facie* rational” to have $C(a_i, \text{ given } F_i) = q$ for some q and all $i \in \mathbb{N}$. Then C is rational only if there exists q such that $C(a_i, \text{ given } F_i) = q$.

I say that a credence function is *prima facie rational* if it is compatible with the probability axioms and accepted principles of rationality, including the sure loss principle (but excluding **ECP**). A *prima facie* rational credence function may be irrational due to some reason that is not conveyed by any accepted principle of rationality (or if **ECP** is accepted as a principle of rationality, because it conflicts with **ECP**).

ECP resembles the principle of indifference, but I am not aware of a derivation on the basis of the principle of indifference. A finite version of an equal credence principle, on the other hand, can be derived from the principle of indifference (Schwarz, 2014). Nevertheless, **ECP** may have the same sort of intuitive appeal as the principle of indifference.

A subjective Bayesian might object to **ECP** because it could be possible to have non-evidential *a priori* information that would require one to have unequal credences in a given situation. This is information that is not encoded in any evidence but is something an agent just ‘knows’. (For example, someone could have *a priori* knowledge implying that near the end of time, coin tosses have a bias towards landing heads, while near the

Footnote 21 continued

know that its frequency converges to the time-ordered limiting frequency. Hence, from the perspective of the agent, there is no sure loss.

²² Williamson’s equal evidence postulate can be interpreted differently from the versions discussed here — a sort of hybrid between the two—but similarly to the liberal equal evidence postulate, this version of the postulate is false. The postulate is formulated by Williamson as $P_{\mathcal{E}}(\theta_i) = q$, where θ_i is the outcome of a chance experiment of type F , $P_{\mathcal{E}}$ is a rational agent’s credence function with respect to the evidence \mathcal{E} , and \mathcal{E} is the evidence that is most pertinent to θ_i . *Most pertinent* means that \mathcal{E} is the smallest set of evidence such that $P_{\mathcal{E}}(\theta_i) = P_{\mathcal{E}'}(\theta_i)$, where \mathcal{E}' is the agent’s complete set of evidence. Williamson assumes that $\mathcal{E} = \{ch(\theta_i) = x\}$, i.e., the most pertinent evidence for θ_i consists only of the single case chance of θ_i (which, on Williamson’s account, implies $f_{\infty}^F(\theta) = x$). This assumption can be held even if one allows that the agent’s full set of evidence \mathcal{E}' for θ_i includes F_1, \dots, F_i , such that the postulate can be usefully applied to a time-ordered sequence of F -experiments. But in that case, the assumption is precisely what’s at stake. The evidence that F -experiments have been conducted $i - 1$ times before might cause a rational agent to choose a different probability for θ_i , in which case \mathcal{E} is *not* the most pertinent evidence. Hence, this version of the equal evidence postulate is false.

beginning of time, they are fair.) I am skeptical of the possibility of *a priori* information that would require a rational agent to have unequal credences—but in any case, the objection can be incorporated by defining ‘*prima facie* rational’ in a way that such *a priori* knowledge has the same status as axioms and principles. (This does have the implication, however, that different rational agents might disagree on whether the conditions of **ECP** are met, although their evidence is identical. Subjectivists would probably be fine with this, but it might make the principle less useful in some respects.)

Another objection contends that in cases of very little information, there is little to say either in favor of equal credences or against equal credences. A counterexample would not exist against either choice. In such a case, it might be reasonable to allow that both equal credences and unequal credences are rational. However, note that to have unequal credences one needs to decide whether credences rise or fall when the experiment is repeated. With very little information, one has nothing to base such a decision on. A principle of indifference of sorts would point towards equal credences in such a situation.

Finally, an objection to **ECP** could be that even when unequal credences are irrational, it might still be rationally permissible to have no credences at all, or to have imprecise credences (as defended by Joyce 2010). Whether an infinite frequency principle can be defended in a situation in which non-existing or imprecise credences are *prima facie* rational could be a topic for future research.

ECP suggests that we accept the Equal Credence Infinite Frequency Principle, according to which credences should be calibrated to the limiting frequency if it is *prima facie* rational to have equal credences in repetitions of the experiment.

EIFP: Let F and a be such that it is *prima facie* rational, for some q and all $i > 0$, to have $C(a_i, \text{ given } F_i) = q$. Let E be any admissible evidence. Then C is rational only if we have $C(a_i, \text{ given } F_i \text{ and } f_\infty^F(a) = q \text{ and } E) = q$.

If **ECP** is correct, then **EIFP** follows by the consistency argument of Sect. 4.2. Hence, it plausibly satisfies the rationality desideratum. Since, unlike **CIFP**, **EIFP** does not rely on a metaphysical primitive of chanciness, it satisfies the precision desideratum.

EIFP is informative—i.e., it constrains beyond consistency—because there are many situations in which both equal and unequal credences are *prima facie* rational (and therefore consistent). **EIFP** scraps a number of these consistent unequal credence functions. It is also informative in the sense of being easily applied. For any given experiment F , one simply checks whether equal credences conflict with any other accepted principle. If not, it is *prima facie* rational to have equal credences, and by **EIFP**, one ought to have equal credences.

The **EIFP** is even more informative than **CIFP**, as there are experiments for which one can know the limiting frequency, and for which equal credences are *prima facie* rational, but which would be denied to be chancy by many philosophers of chance. Take the following example. Let F describe a computer program whose output x_i is the next number from either of two sequences, $(a_i) = (1, 0, 1, 0, \dots)$ or $(b_i) = (0, 1, 0, 1, \dots)$. Each time it is run, it chooses a sequence on the basis of the current time: if the current second is even, it takes the next number from (a_i) ; if the current

second is odd, from (b_i) . It is clear that regardless of how this experiment is repeated, the limiting frequency of output '1' is always $1/2$.

Consider your credence in output '1', supposing you only know that F is run, but not what the current time is. This credence could be written as

$$C(x_i = 1, \text{ given } F_i \text{ and } f_{\infty}^F(1) = 1/2).$$

There is obviously no reason of rationality that requires you to have different credences for repetitions of F . Hence, by **EIFP**, we have

$$C(x_i = 1, \text{ given } F_i \text{ and } f_{\infty}^F(1) = 1/2) = 1/2,$$

which seems to be the intuitively correct credence.

However, some would deny that F is chancy. To see why, consider what a sequence of its outputs looks like. When there are two ones in a row, these two ones must have come from different sequences. After two ones, the next number to be drawn from either sequence is a zero. Hence, whenever there are two ones in a row, the next output must be a zero; and when there are two zeros in a row, the next output must be a one. Such a sequence is far from random. According to some, however, chancy experiments are thought to yield random sequences most of the time (Eagle, 2016). If that is the case, F cannot be chancy.

7 Conclusion

I argued that available infinite frequency principles all suffer from problems, but I introduced a new principle that does not seem to suffer from these problems. Future theorists of rational credence should take a few things into account.

First, the Naive Infinite Frequency Principle does not satisfy the rationality desideratum and should be avoided. Relatedly, arguments for frequency principles work only if you include the assumption that different instantiations of the experiment have identical credences.

Second, the Chancy Infinite Frequency Principle suffers from a variety of problems depending on how it is interpreted. If chanciness is left undefined, it does not satisfy the precision desideratum. If it is defined using one of the existing definitions of chanciness, it does not satisfy either the informativeness desideratum or the rationality desideratum. Further research is required to investigate whether chanciness can be given a satisfactory definition at all.

Third, given the problems with the Chancy Infinite Frequency Principle, theorists should reconsider whether they want to depend on it. The Equal Credence Infinite Frequency Principle does the same job, satisfies all three desiderata, and is more informative than the Chancy Infinite Frequency Principle.

7.1 A coda: back to chanciness

As shown in example 3 in Sect. 2.3, a frequency principle can be applied to test a theory of chance. Similarly, one can use it to define a theory of chance in such a way that it passes this test. I will give some suggestions for using an infinite frequency principle to gain insight into the nature of chance. I consider three possible definitions of chance similar to Mellor's definition in *The Facts of Causation* (1995).

First, Mellor argues that chances are properties that meet the Necessity, Evidence, and Frequency conditions. While Mellor's theory is usually considered to describe single-case chance, it is applied here to type chances for consistency with the rest of the paper. An outcome a of an experiment F has the chance $ch_F(a)$ if and only if the following conditions hold.

- (1) If $ch_F(a) = 1$ and F_i are a fact, then a_i is a fact. (Necessity)
- (2) For all rational credence functions C , we have $C(a_i, \text{given } F_i \text{ and } ch_F(a) = p) = p$. (Evidence)
- (3) We have $f_\infty^F(a) = ch_F(a)$. (Frequency)

Condition (2) ensures that chances satisfy the principal principle by definition. Condition (3) further constrains chances to situations in which the limiting frequency is defined.

One might worry that (2) is overly broad, or that (2) is uninformative as it requires detailed knowledge of what is rational. See Strevens (1999) for a criticism of the latter type.²³

Second, we can replace (2) with the following, less demanding condition.

- (2') It is *prima facie* rational, for some q and all $i > 0$, to have $C(a_i, \text{given } F_i) = q$.

Suppose that a theory \mathbf{T} defines chances as properties that meet (1), (2'), and (3). (Note that (2') is still a condition on chance, since it constrains which experiments F can be assigned a chance, i.e., which experiments F are chancy.) Consider that (2') in combination with (3) implies that for every \mathbf{T} -chance, the conditions for invoking **EIFP** are met, with $f_\infty^F(a) = ch_F(a)$. Hence, if **EIFP** is true as a principle of rationality, then it is rational to set one's credence to the chance for every \mathbf{T} -chance. In other words, \mathbf{T} -chances automatically satisfy the principal principle.

This definition of chanciness would lead us to consider experiments chancy that many believe not to be chancy. See the last example of the previous section: repeating that program leads to sequences that are non-random, but chancy experiments are often thought to produce random sequences if they are repeated. A different way of putting it is that repetitions of a chancy experiment should be independent of each other.

A third definition of chance suggests itself that rules out such cases, using exchangeability:

- (2'') For all rational credence functions C , a_1, a_2, \dots given F are exchangeable.

Exchangeability is equivalent to the situation that the outcomes a_1, a_2, \dots are conditionally independent, conditional on the limiting frequency. Hence, (2'') would restrict

²³ Strevens (1999, p. 65) criticizes the practice of "tweaking" the definition of chance such that chances satisfy the principal principle by definition, as this only moves the philosophical problem of justifying a principle of direct inference to problems relating to the existence and inference of objective probabilities.

the phenomenon of chanciness to cases in which repetitions of the same chance experiment are in some sense i.i.d. However, a possible downside of (2'') compared to (2') is that it is quite unclear when rationality requires exchangeability, or even that exchangeability is ever rationally required. If that problem is not resolved, a theory of chance that uses (2'') is uninformative.

Perhaps it is possible to do away with all references to rationality in a definition of chanciness, or perhaps we cannot do better than a definition of the above type. This could be an interesting topic for future research.

Acknowledgements I am grateful to Christopher Clarke, Wolfgang Schwarz, Conrad Heilmann, and the anonymous reviewers for helpful comments and suggestions.

Funding This work has received funding from the European Research Council under the European Union's Horizon 2020 Research and Innovation Programme, under Grant Agreement No 715530.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ackermans, L.B. (forthcoming). Deterministic propensities. In Probability, causation and inequality of opportunity. Rotterdam: Erasmus University Rotterdam. (PhD thesis)
- Albert, M. (2005). Should Bayesians bet where frequentists fear to tread? *Philosophy of Science*, 72(4), 584–593. <https://doi.org/10.1086/505185>
- Berkovitz, J. (2015). The propensity interpretation of probability: A re-evaluation. *Erkenntnis*, 80(3), 629–711. <https://doi.org/10.1007/s10670-014-9716-8>
- Carnap, R. (1962). *Logical foundations of probability* (2nd ed.). The University of Chicago Press.
- Church, A. (1940). On the concept of a random sequence. *Bulletin of the American Mathematical Society*, 46(2), 130–135.
- De Finetti, B. (1937). La prévision: Ses lois logiques, ses sources subjectives. *Annales de l'institut Henri Poincaré*, 7(1), 1–68.
- Eagle, A. (2004). Twenty-one arguments against propensity analyses of probability. *Erkenntnis*, 60(3), 371–416. <https://doi.org/10.1023/B:ERKE.0000023408.61887.6a>
- Eagle, A. (2016). Probability and randomness. In A. Hájek & C. Hitchcock (Eds.), *Oxford handbook of probability and philosophy* (pp. 440–459). <https://doi.org/10.1093/oxfordhb/9780199607617.013.22>
- Fetzer, J. H. (1981). Probability and explanation. *Synthese*, 48, 371–408. <https://doi.org/10.1007/BF01063986>
- Fischer, M. (2021). On the principal principle and imprecise subjective Bayesianism. *European Journal for Philosophy of Science*, 11(2), 47.
- Glynn, L. (2009). Deterministic chance. *The British Journal for the Philosophy of Science*, 61(1), 51–80. <https://doi.org/10.1093/bjps/axp020>
- Hájek, A. (2007). The reference class problem is your problem too. *Synthese*, 156(3), 563–585. <https://doi.org/10.1007/s11229-006-9138-5>
- Hájek, A. (2009a). Dutch book arguments. In P. P. Paul Anand & C. Puppe (Eds.), *The handbook of rational and social choice*. <https://doi.org/10.1093/acprof:oso/9780199290420.003.0008>
- Hájek, A. (2009). Fifteen arguments against hypothetical frequentism. *Erkenntnis*, 70(2), 211–235. <https://doi.org/10.1007/s10670-009-9154-1>

- Hawthorne, J., Landes, J., Wallmann, C., & Williamson, J. (2017). The principal principle implies the principle of indifference. *The British Journal for the Philosophy of Science*, 68(1), 123–131. <https://doi.org/10.1093/bjps/axv030>
- Hofer, C. (2007). the third way on objective probability: A sceptic's guide to objective chance. *Mind*, 116(463), 549–596. <https://doi.org/10.1093/mind/fzm549>
- Howson, C., & Urbach, P. (1993). *Scientific reasoning: The Bayesian approach* (2nd ed.). Open Court. (First edition published in 1989.).
- Howson, C., & Urbach, P. (2006). *Scientific reasoning: The Bayesian approach* (3rd ed.). Open Court. (First edition published in 1989.).
- Joyce, J. M. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives*, 24, 281–323.
- Landes, J., Wallmann, C., & Williamson, J. (2021). The principal principle, admissibility, and normal informal standards of what is reasonable. *European Journal for Philosophy of Science*, 11(2), 1–15. <https://doi.org/10.1007/s13194-021-00352-x>
- Lewis, D. (1980). A subjectivist's guide to objective chance. In R. C. Jeffrey (Ed.), *Studies in inductive logic and probability* (Vol. 2, pp. 263–293). University of California Press.
- Mellor, D. H. (1995). *The facts of causation*. *International Library of Philosophy*. Routledge.
- Mellor, D. H. (2005). *Probability: A philosophical introduction*. Routledge.
- Niiniluoto, I. (1981). Statistical explanation reconsidered. *Synthese*, 48(3), 437–472. <https://doi.org/10.1007/BF01063988>
- Pivato, M., & List, C. (2015). Emergent chance. *The Philosophical Review*, 124(1), 119–152. <https://doi.org/10.1215/00318108-2812670>
- Ramsey, F. P. (1931). Truth and probability. In R. B. Braithwaite (Ed.), *Foundations of mathematics and other essays* (pp. 156–198). Routledge & Kegan Paul.
- Schaffer, J. (2003). Principled chances. *The British Journal for the Philosophy of Science*, 54(1), 27–41. <https://doi.org/10.1093/bjps/54.1.27>
- Schaffer, J. (2007). Deterministic chance? *The British Journal for the Philosophy of Science*, 58(2), 113–140. <https://doi.org/10.1093/bjps/axm002>
- Schervish, M. J. (1995). *Theory of statistics*. Springer series in statistics Springer.
- Schwarz, W. (2014). Proving the principal principle. In A. Wilson (Ed.), *Chance and temporal asymmetry*. <https://doi.org/10.1093/acprof:oso/9780199673421.003.0004>
- Sober, E. (2010). Evolutionary theory and the reality of macro-probabilities. In E. Eells & J. Fetzer (Eds.), *The place of probability in science: In honor of ellery eells (1953-2006)* (pp. 133-161). https://doi.org/10.1007/978-90-481-3615-5_6
- Strevens, M. (1999). Objective probability as a guide to the world. *Philosophical Studies*, 95(3), 243–275. <https://doi.org/10.1023/A:1004256510606>
- Suppes, P. (1987). Propensity representations of probability. *Erkenntnis*, 26(3), 335–358. <https://doi.org/10.1007/BF00167720>
- Thorn, P. D. (2012). Two problems of direct inference. *Erkenntnis*, 76(3), 299–318. <https://doi.org/10.1007/s10670-011-9319-6>
- Thorn, P. D. (2017). On the preference for more specific reference classes. *Synthese*, 194(6), 2025–2051. <https://doi.org/10.1007/s11229-016-1035-y>
- von Mises, R. (1957). *Probability, statistics, and truth* (2nd ed.). George Allen & Unwin Ltd.
- Wallmann, C. (2017). A bayesian solution to the conflict of narrowness and precision in direct inference. *Journal for General Philosophy of Science*, 48(3), 485–500. <https://doi.org/10.1007/s10838-017-9368-x>
- Wallmann, C., & Hawthorne, J. (2020). Admissibility troubles for Bayesian direct inference principles. *Erkenntnis*, 85(4), 957–993. <https://doi.org/10.1007/s10670-018-0070-0>
- Wallmann, C., & Williamson, J. (2020). The principal principle and subjective Bayesianism. *European Journal for Philosophy of Science*, 10(3), 3. <https://doi.org/10.1007/s13194-019-0266-4>
- Williamson, J. (2010). *In defence of objective Bayesianism*. Oxford University Press.