Conditional Degree of Belief and Bayesian Inference

Jan Sprenger*†

Why are conditional degrees of belief in an observation E, given a statistical hypothesis H, aligned with the objective probabilities expressed by H? After showing that standard replies (ratio analysis of conditional probability, chance-credence coordination) are not satisfactory, I develop a suppositional analysis of conditional degree of belief, transferring Ramsey's classical proposal to statistical inference. The analysis saves the alignment, explains the role of chance-credence coordination, and rebuts the charge of arbitrary assessment of evidence in Bayesian inference. Finally, I explore the implications of this analysis for Bayesian reasoning with idealized models in science.

1. Introduction. Bayesian inference is a well-established theory of uncertain reasoning that represents agents' epistemic attitudes—their degrees of belief—by the laws of probability (e.g., Jeffrey 1965; de Finetti 1972; Savage 1972; Earman 1992; Bovens and Hartmann 2003). A probability function $p(\mathrm{H})$ represents a rational agent's degree of belief that H is true. On learning evidence E, the agent adopts a posterior belief in H according to the rule of conditionalization: $p^{\mathrm{E}}(\mathrm{H}) = p(\mathrm{H}|\mathrm{E})$. Such posterior degrees of belief serve as a basis for assessing hypotheses and making decisions—also in the context of public policy. For example, the assessment reports of the

Received October 2018; revised June 2019.

*To contact the author, please write to: Center for Logic, Language and Cognition, Department of Philosophy and Educational Science, Università degli Studi di Torino, Via Sant'Ottavio 20, 10124 Torino, Italy; e-mail: jan.sprenger@unito.it.

†I would like to thank Guido Bacciagaluppi, Max Bialek, Claus Beisbart, Colin Elliot, Alan Hájek, Stephan Hartmann, Jan-Willem Romeijn, Carlotta Pavese, Tom Sterkenburg, Olav Vassend, and audiences in Groningen, Sestri Levante, Tilburg, and Turin for their valuable feedback. Furthermore, three anonymous referees of *Philosophy of Science* contributed to improving this article. Research on this article was supported through the Starting Investigator grant 640638, "Making Scientific Inferences More Objective," by the European Research Council.

Philosophy of Science, 87 (April 2020) pp. 319–335. 0031-8248/2020/8702-0006$10.00
Copyright 2020 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/707554 Published online by Cambridge University Press

International Panel for Climatic Change evaluate the probability of future events as experts' subjective degrees of belief.

Why are these posterior degrees of belief something else than arbitrary subjective attitudes? Why can they guide rational and efficient decisions? Presumably it is because they are in some way informed by objective evidence. Indeed, if we look at Bayes's theorem,

$$p(\mathrm{H}|\mathrm{E}) = p(\mathrm{H})\frac{p(\mathrm{E}|\mathrm{H})}{p(\mathrm{E})} = \left(1 + \frac{1 - p(\mathrm{H})}{p(\mathrm{H})} \cdot \frac{p(\mathrm{E}|\neg\mathrm{H})}{p(\mathrm{E}|\mathrm{H})}\right)^{-1},$$

we see that an agent's posterior degree of belief in H depends on three factors: her prior degree of belief in H, $p(\mathrm{H})$, and the conditional degrees of belief $p(\mathrm{E}|\mathrm{H})$ and $p(\mathrm{E}|\neg\mathrm{H})$—often called the likelihoods of H and ¬H on E. Bayesians contend that as long as the prior degrees of belief in H are not too extreme, a "well-designed experiment . . . will swamp divergent prior distributions with the clarity and sharpness of its results, and thereby render insignificant the diversity of prior opinion" (Suppes 1966, 204). Clearly, any such merging-of-opinion argument makes the tacit assumption that $p(\mathrm{E}|\mathrm{H})$ and $p(\mathrm{E}|\neg\mathrm{H})$ are objectively constrained.

The reliance of Bayesians on such constraints is even more explicit in the Bayes factor—a standard measure for summarizing experimental observations and quantifying the weight of evidence in favor of a hypothesis (Jeffreys 1961; Kass and Raftery 1995). It is defined as the ratio of prior and posterior odds between two competing hypotheses H_0 and H_1:

$$\mathrm{BF}_{10}(\mathrm{E}) = \frac{p(\mathrm{H}_1|\mathrm{E})/p(\mathrm{H}_0|\mathrm{E})}{p(\mathrm{H}_1)/p(\mathrm{H}_0)}.$$

It follows from Bayes's theorem that the Bayes factor is a ratio of two conditional degrees of belief: $\mathrm{BF}_{10}(\mathrm{E}) = p(\mathrm{E}|\mathrm{H}_1)/p(\mathrm{E}|\mathrm{H}_0)$. Thus, the Bayes factor is only as objective and nonarbitrary as $p(\mathrm{E}|\mathrm{H}_0)$ and $p(\mathrm{E}|\mathrm{H}_1)$ are.[1]

The idea that such conditional degrees of belief are rationally constrained and temper the influence of subjectively chosen priors stands at the basis of many attempts to defend the rationality and objectivity of Bayesian inference (e.g., Earman 1992, chap. 6). It is usually taken for granted that they are aligned with objective probabilities derived from the relevant statistical model. For example, if H denotes the hypothesis that a die is fair and E the outcome

1. This claim presupposes that H_0 and H_1 are two precise point hypotheses—an assumption made throughout the article for reasons of simplicity. Section 5 briefly discusses the general case.

https://doi.org/10.1086/707554 Published online by Cambridge University Press

of two sixes in two tosses, then it appears rational to have the conditional degree of belief $p(E|H) = 1/36$. However, none of the familiar stories for justifying this alignment (e.g., reliance on chance-credence coordination principles) is convincing as it stands.

Pointing out this justification gap is the main task of the negative part of the article. In particular, I argue that neither the ratio analysis of conditional probability nor chance-credence coordination principles explain why conditional degrees of belief are rationally constrained by the corresponding objective probabilities (sec. 2). The constructive part of the article solves the problem in a Ramsey–de Finetti spirit: $p(E|H)$ is the degree of belief in the occurrence of E upon supposing that the target system's behavior is described by H. I work out the details of this suppositional analysis in the context of statistical inference: the relevant set of possible worlds for evaluating conditional degrees of belief, their alignment with density functions of statistical models, and how chance-credence coordination guides Bayesian inference and supports claims to objectivity (sec. 3).

In the final part of the article, I explore the general implications of my approach, focusing on a pertinent problem of Bayesian inference: the interpretation of highly idealized statistical models in which important causal factors are omitted or functional dependencies are simplified. Such models are used in disciplines as diverse as psychology, economics, and climate science. In such cases, it would be inappropriate to interpret the probability of a model as the degree of belief in its (approximate) truth. Nonetheless, as Bayesian reasoners, we rank different idealized models according to their posterior probability, and we use these rankings in inference and decision making. So we need to explain what these probabilities mean, if not degrees of belief in the truth of the model.

I argue that the problem vanishes when all probabilities in Bayesian inference, including prior and posterior degrees of belief, are understood as conditional degrees of belief relative to an overarching model. Then I explain how this extension of the suppositional analysis squares with various principles for determining rational prior degrees of belief, including a recent proposal by Olav Vassend (sec. 4). Moreover I show how we can use Bayesian models for prediction, theory evaluation, and decision making, even when models are highly idealized and not faithful to reality (sec. 5). Finally I wrap up the results of the article (sec. 6).

The suppositional approach is not novel. Ramsey (1929/1990) famously argued that a conditional degree of belief in a proposition E given another proposition H is determined by supposing H and reasoning on that basis about E. However, my article is, to the best of my knowledge, the first one to explicate the mechanics of the suppositional approach in the context of statistical inference, to make precise the role of chance-credence coordination in this process, and to explain why such conditional degrees of belief

https://doi.org/10.1086/707554 Published online by Cambridge University Press

are universally shared.[2] It is also the first exploration of the implications of the suppositional approach for contexts in which no model is a serious contender for (approximative) truth and for practical decisions based on statistical models. Thus, it provides the conceptual groundwork for numerous applications of Bayesian inference in statistics and other domains of science.

2. "The Equality" and Probability in Statistical Models. What constrains the conditional degree of belief in an observation E, given a statistical hypothesis H? A classical illustration is an inference about the bias of a coin. The hypotheses (H_μ, $\mu \in (0, 1)$) describe how likely it is that the coin will come up heads on any individual toss. When the tosses are independent and identically distributed (i.i.d.), we can describe the outcome of N repeated coin tosses by the observation E_k ($=k$ heads and $N - k$ tails), whose probability follows the binomial probability density function $\rho_{H_\mu}(E_k) = \binom{N}{k} \mu^k (1 - \mu)^{N-k}$.[3]

Suppose that we consider the hypothesis that a coin is slightly biased toward tails: it comes up heads only 40% of the time ($H_0 : \mu = .4$). This implies that the probability of observing two heads in two i.i.d. tosses ($=E$) is equal to $\rho_{H_0}(E) = \binom{2}{2} (.4)^2 (.6)^0 = .16$. Bayesian reasoners align their conditional degrees of belief $\rho(E|H_0)$ with the relevant value of the probability density function ρ_{H_0}, that is, $\rho_{H_0}(E)$ (e.g., Bernardo and Smith 1994; Howson and Urbach 2006). And since the latter is uniquely determined, so is the former. Having "objective" conditional degrees of belief $p(E|H_0)$ leads to (approximate) long-run consensus on posterior distributions and unanimous assessments of the strength of observed evidence, for example, via the Bayes factor. For the above evidence E and the two competing hypotheses $H_0 : \mu = .4$ and $H_1 : \mu = .8$, all of us will presumably adopt the conditional degrees of belief $p(E|H_0) = .4 \times .4 = .16$ and $p(E|H_1) = .9 \times .9 = .81$ and report a Bayes factor $BF_{10}(E) = .81/.16 \approx 5.06$, corresponding to moderate evidence for H_1. This evidential judgment is supposed to be shared by all Bayesian reasoners, regardless of their priors over H_0 and H_1. It rests, however, on an alignment between density functions of a statistical model and conditional degrees of belief that is not easy to justify.

2. The closest relative is perhaps Levi (1980, esp. chap. 12), but Levi's conceptual framework is very different, from the central role of confirmational commitments in Bayesian inference to his handling of chance predicates and the absence of possible-world semantics for spelling out conditional degree of belief. Moreover, Levi groups both ontic and statistical probability under the label of (objective) chance.

3. The binomial distribution describes the expected number of successes in a sequence of i.i.d. Bernoulli (i.e., success-or-failure) trials. Together with the sample space $\mathcal{S} = \{H, T\}^N$, the probability distributions ρ_{H_μ} over $\mathcal{S}$ constitute a statistical model.

https://doi.org/10.1086/707554 Published online by Cambridge University Press

In other words, we require a satisfactory answer to the following question:

> **Main Question.** What justifies the equality between conditional degrees of belief and the corresponding probability densities?
> (THE EQUALITY) $p(\mathrm{E}|\mathrm{H}) = \rho_{\mathrm{H}}(\mathrm{E})$.

A traditional approach to conditional degrees of belief, proposed by various textbooks on Bayesian inference (e.g., Jackson 1991; Earman 1992; Skyrms 2000; Howson and Urbach 2006), evaluates them as the ratio of two unconditional probabilities: $p(\mathrm{E}|\mathrm{H}) = p(\mathrm{E} \wedge \mathrm{H})/p(\mathrm{H})$ whenever $p(\mathrm{H}) > 0$. While this ratio analysis (Hájek 2003) is uncontroversial as a mathematical constraint on conditional probability, it does not explain THE EQUALITY. First, if $\rho_{\mathrm{H}}(\mathrm{E})$ determines $p(\mathrm{E}|\mathrm{H})$, it must do so directly and not be mediated by unconditional degrees of belief. The counterparts of $p(\mathrm{E} \wedge \mathrm{H})$ and $p(\mathrm{H})$ are undefined within the statistical model (e.g., the binomial distribution). Second, ratio analysis neglects a robust empirical fact: we usually evaluate conditional degrees of belief $p(\mathrm{E}|\mathrm{H})$ directly, rather than by reasoning about $p(\mathrm{E} \wedge \mathrm{H})$ and $p(\mathrm{H})$ and calculating their ratio (see also Hájek [2003] and the experiments in Zhao, Shah, and Osherson [2009]). Third and last, ratio analysis is silent whenever $p(\mathrm{H}) = 0$, but such hypotheses are omnipresent in statistical inference with real-valued parameters. For example, a uniform prior distribution—in fact, any continuous distribution—over the parameter μ in the binomial model implies that any precise hypothesis such as $\mathrm{H} : \mu = .4$ has probability zero.

A frequently used alternative strategy for justifying THE EQUALITY consists in invoking a chance-credence coordination principle: subjective credences should follow known objective chances. The most famous of these is the Principal Principle (PP; Lewis 1980): the initial credence function of a rational agent, conditional on the proposition that the physical chance of E takes value x, should be equal to x. A similar intuition with an eye on applications to statistical inference is expressed by the Principle of Direct Inference (PDI; e.g., Reichenbach 1949; Kyburg 1974; Levi 1977, 1980): for instance, if I know that a die is unbiased, I should assign degree of belief 1/6 that any particular number will come up.

Transferring these principles to THE EQUALITY is, however, not as straightforward as it looks. True, the value $\rho_{\mathrm{H}}(\mathrm{E})$ does not depend on subjective credences but on the objective properties of a given statistical model. This seems to be a sufficient reason for classifying $\rho_{\mathrm{H}}(\mathrm{E})$ as an objective chance, and then PP or PDI determines the value of the conditional degree of belief $p(\mathrm{E}|\mathrm{H})$ (see also Earman 1992, 54–56).

This strategy neglects that chance-credence coordination principles understand objective chances as making empirical statements: their values depend on "facts entailed by the overall pattern of events and processes in the *actual* world" (Hoefer 2007, 549). Dependent on the preferred conception

https://doi.org/10.1086/707554 Published online by Cambridge University Press

of objective chance, such facts could be the setup of a statistical experiment or the composition and precise shape of a die that we roll. However, the truth conditions of sentences such as $\rho_H(E) = 1/36$ are entirely internal to the statistical model. Suppose that H denotes the hypothesis that a die in front of us is unbiased, and E denotes the observation of two sixes in two i.i.d. rolls. Then the sentence

> (Dice Roll) When we roll an unbiased die twice, the chance of observing two sixes is 1/36.

has no empirical content—it may even strike us as analytically true. Note that Dice Roll does not refer to real-world properties or events: even if no perfectly unbiased dice existed in the actual world (and perhaps this is actually the case), the sentence would still be true. The probability $\rho_H(E)$ is objective in the sense of being subject independent but is not a physical chance in the sense of being realized in the actual world (Rosenthal 2004; Sprenger 2010). This diagnosis is typical of probability in statistical models.

For this reason, neither the PP nor the PDI solves our problem. The principles coordinate our degrees of belief with known chancy properties of the actual world. But the objective probabilities in question, $\rho_H(E)$, do not express physically realized chances. Therefore, standard chance-credence coordination principles cannot directly justify THE EQUALITY. A more sophisticated story has to be told.

3. The Suppositional Analysis. This section develops and defends a suppositional analysis of conditional degree of belief: conditional degrees of belief constitute a primitive epistemic concept and we determine our degree of belief in E given H by supposing that H is true. Two famous Bayesians—the British philosopher Frank P. Ramsey and the Italian statistician Bruno de Finetti (1972, 2008)—have proposed this view in the literature. I focus on Ramsey since de Finetti also requires that H be a verifiable event if $p(E|H)$ is to be meaningful (de Finetti 1972, 193). This verificationism is unnecessarily restrictive for our purposes.

Here is Ramsey's famous analysis of conditional degrees of belief: "If two people are arguing 'if H will E?' and both are in doubt as to H, they are adding H hypothetically to their stock of knowledge and arguing on that basis about E. We can say that they are *fixing their degrees of belief in* E *given* H" (1929/1990, 155, my emphasis). Put differently, we evaluate the conditional degree of belief $p(E|H)$ by supposing the truth of the conditioning proposition H and by assessing the plausibility of E given this supposition. Ramsey's analysis has also inspired various accounts of evaluating (the probability of) indicative conditionals (e.g., Stalnaker 1968; Adams 1975; Levi 1996), but these questions go beyond the scope of this article.

https://doi.org/10.1086/707554 Published online by Cambridge University Press

While we have an intuitive grasp of how Ramsey's proposal is supposed to work, we need a more detailed account of its mechanics to explain why statistical reasoners typically agree on the relevant conditional degrees of belief. Consider a target system S (e.g., repeated dice rolls) described by a statistical hypothesis H (e.g., "the die is fair"). Supposing H defines a possible world ω_H, or more precisely a set of possible worlds, where the behavior of S is governed by ρ_H. Given that these possible worlds may differ from one another in features that are unrelated to S, which one is relevant for fixing our degrees of belief? Do we need to choose the closest possible world—a notoriously vague and difficult concept (Lewis 1973a)—to settle the matter?

Fortunately, choosing is not necessary. Let $[\omega_{H,S}] \subset \mathcal{W}$ denote the set of worlds where the behavior of S is governed by the probability law H. Supposing H is best explicated as restricting the space of relevant possible worlds to $[\omega_{H,S}]$. In particular, in any such world, the objective chance of an observation E is given by $\rho_H(E)$.[4]

The differences between the elements of $[\omega_{H,S}]$ are not relevant for our purposes. Typically, the scope of a statistical model does not go beyond the target system it aims to model. For example, in an experiment in which we roll a die, the hypotheses correspond to (multinomial) distributions describing the die's specific properties. Similarly, the possible outcomes E, E′, E″, and so on (e.g., three sixes in a row), are contained in our statistical model of S. In any possible world that belongs to $[\omega_{H,S}]$, the outcome E will therefore have the same probability. This invariance is a notable difference between applying the suppositional analysis in the context of statistical inference and to conditional degrees of belief more generally.

Supposing H may be in conflict with available background knowledge about the target system. For this reason, my interpretation of conditional degrees of belief differs from Ramsey's in a crucial nuance: where Ramsey suggested that H is added to existing background knowledge, on my account H may also overrule conflicting information. In such cases, we obtain a genuinely counterfactual interpretation of conditional degrees of belief. This is often necessary: we may know that a given die is biased, that the rolls are not i.i.d., and so on. But for my analysis, it does not matter whether assuming H is consistent or in conflict with our background knowledge: the above recipe for constructing the set $[\omega_{H,S}]$ applies in either case.

The dice-rolling example shows how the suppositional analysis ensures alignment of conditional degrees of belief with probability densities. Let H denote the hypothesis that the die on the table is fair. Consider a world $\omega_H \in [\omega_{H,S}]$.

4. I would like to thank an anonymous reviewer of this journal for suggesting this simple definition. Originally, the relevant class of possible worlds was defined via an equivalence relation on possible worlds (=assigning the same probability law to S), but that approach would be unnecessarily technical.

https://doi.org/10.1086/707554 Published online by Cambridge University Press

As explained above, supposing H implies that within ω_H, the objective, physical chance of rolling (at least) one six in one toss is 1/6, in two (i.i.d.) tosses it is 11/36, and so on. Since ω_H is by definition a chancy world, these chances should inform our degrees of belief via the PDI or the PP. After all, ρ_H describes the physical chances that hold for S in ω_H, and we have no reason to challenge the rationality of PDI or PP in this case.[5] Thus, for any event E in target system S, our (unconditional) degrees of belief within ω_H should satisfy $p_{\omega_H}(E) = \rho_H(E)$. By definition of the suppositional analysis of conditional degrees of belief, we also have $p(E|H) = p_{\omega_H}(E)$. Combining both equations yields $p(E|H) = \rho_H(E)$: conditional degrees of belief track probability density functions (see also fig. 1).[6] Taken together with uncontroversial principles for chance-credence coordination, the suppositional analysis of conditional degrees of belief establishes THE EQUALITY and explains the seemingly analytic character of sentences such as Dice Roll. In particular, there is no room for rational disagreement on such conditional degrees of belief.

This agreement transfers to statistical measures of evidence that are derived from these conditional degrees of belief and play a pivotal role in statistical inference, such as Bayes factors. Hence, we can explain where the objective elements in Bayesian inference come from: probability density functions determine conditional degrees of belief (i.e., the likelihoods) and constrain measures of evidence such as Bayes factors. Via Bayes's theorem, they also constrain posterior probabilities (see sec. 5 for more details).

It is important to understand the role of the PDI and the PP in the suppositional analysis. Standardly, both principles apply to real-world, ontic chances, for example, "the chance of this atom decaying in the next hour is 1/3" or "the chance of a zero in the spin of this roulette wheel is 1/37." The principles claim that degrees of belief should mirror physical chances whenever we know them. Compare this to the picture that we sketch for conditional degrees of belief: we do not deal with real-world chances; rather we observe that in the worlds in the set $[\omega_{H,S}]$, the physical chance of E is given by $\rho_H(E)$. In other words, we do not apply PDI or PP in the actual world $\omega_@$ but in a counterfactual world ω_H described by H. By supposing that the occurrence of E is genuinely chancy and follows the probability law $\rho_H(E)$, the suppositional analysis gives a role to chance-credence coordination principles in statistical inference, explaining why our conditional degree of belief in E given H is

5. Some of our actual background information could make an application of PP *inadmissible* in the sense of Lewis (1980). I have not found a convincing example myself, but if somebody did, this worry could be addressed by restricting $[\omega_{H,S}]$ to a subset where no background assumptions interfere with an application of PP. Spelling this strategy out in detail is an exciting topic for future research.

6. Implicitly, this argument may require the assumption of *conglomerability* (e.g., Dubins 1975): if $H = \cup H_i$ for disjoint H_i and $p(E|H_i) = x$ for all indices i, then also $p(E|H) = x$.

https://doi.org/10.1086/707554 Published online by Cambridge University Press

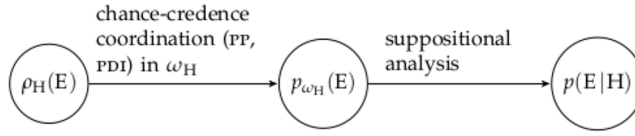


Figure 1. Two steps in justifying THE EQUALITY: chance-credence coordination transfers the probability density $\rho_H(E)$ to the rational ω_H degree of belief in E, $p_{\omega_H}(E)$, and the suppositional analysis connects that value to $p(E|H)$.

uniquely determined and obeys THE EQUALITY. Note that our application of PDI or PP pertains to unconditional degrees of belief (i.e., $p_{\omega_H}(E)$) and is therefore fully analogous to those physical-chance examples that motivate the principles in the first place.

Transferring chance-credence coordination from actual to counterfactual worlds is a distinct strength of the proposed account, and a crucial difference to competing accounts that introduce a distinct chance predicate, and chance propositions in our corpus of background knowledge (Levi 1980, 254–56). So much the more as the existence of physical chances in the real world is a contested issue, especially outside some foundational areas of science such as quantum physics. If chance-credence coordination in the actual world were supposed to justify THE EQUALITY, the Bayesian would need an argument that physical chances actually exist for the system she is studying and that they are expressed by the values of ρ_H. Such a claim would be hard to prove and make strong ontological commitments. Instead, the probabilities given by the model's density functions ρ_H should be understood as physical chances in hypothetical scenarios. Chance-credence coordination should apply upon the supposition of such a hypothetical scenario and not in the actual world.

This agnostic attitude to physical chance matches the practice of non-Bayesian statistical inference, too. Here are the thoughts of the great frequentist statistician Ronald A. Fisher on conditional probability in hypothesis testing: "In general tests of significance are based on *hypothetical* probabilities calculated from their null hypotheses. They *do not lead to any probability statements about the real world*" (1956, 44). That is, Fisher is emphatic that the conditional probability of data given some hypothesis has a hypothetical character and is not a physically realized objective chance. Probabilities are useful instruments of inference, not components of the actual world. According to Fisher, probabilistic reasoning and hypothesis testing is essentially counterfactual: it is based on the probability of observations under an idealized and most likely false hypothesis that we suppose for the sake of the argument.

Before proceeding, I recap the essential elements of the proposed suppositional analysis. The starting point is Ramsey: the conditional degree of belief in observation E given statistical hypothesis H is equal to the belief in E that we have upon supposing that H is the true model of target system S. This hypothetical

https://doi.org/10.1086/707554 Published online by Cambridge University Press

scenario corresponds to a set of possible worlds $[\omega_{H,S}]$ where S follows the probabilistic law defined by H. In this stochastic world $[\omega_{H,S}]$, we apply standard chance-credence coordination principles and calibrate our degree of belief in E with the known objective chance of E, given by the density function $\rho_H(E)$. Thus, we obtain THE EQUALITY. The suppositional approach to conditional degree of belief also relieves us of the worry to make sense of $p(E|H)$ in the frequently occurring cases in which $p(H) = 0$. In these cases, probabilities of the type $p(E|H)$ fall outside the scope of (standard) ratio analysis, but for us, they have determinate values since the suppositional analysis is feasible regardless of the probability of H. I will now explore the implications of the suppositional analysis for Bayesian inference in general.

4. Implications: Model-Relative Prior Probabilities. The suppositional analysis has focused on the relationship between statistical hypothesis H and observation E and said little about the background assumptions in the statistical model $\mathcal{M}$. Consider again the case of tossing a coin, assuming that the tosses are i.i.d. This assumption is not expressed in a hypothesis such as "the coin is fair" or "the coin is slightly biased toward tails." Differences between the competing hypotheses are typically expressed by means of different parameter values, such as $H : \mu = 1/2$ versus $H' : \mu = .4$, $H'' : \mu = .9$, and so on. The statement of the hypothesis does not comprise essential assumptions about the experimental setup such as independence and identical distribution of the coin tosses.

Introducing the notion of a statistical model $\mathcal{M} = (\mathcal{S}; \mathcal{P})$ (Cox and Hinkley 1974; McCullagh 2002) helps to clarify matters. It consists of two parts: the sample space $\mathcal{S}$ and the set of probability distributions $\mathcal{P}$ over $\mathcal{S}$, such as the binomial distributions $B(N, \mu)$ that describe the repeated toss of a coin. All these distributions presuppose the sample size N and that the individual tosses are i.i.d. Therefore, the hypotheses in $\mathcal{M}$ only differ regarding the value of μ, and for purposes of comparisons we can write them as $H : \mu = 1/2$, $H' : \mu = .4$, and so on.

Note that real coin tosses are often far from being i.i.d. Nevertheless, restriction to the family of binomial distributions can be a useful and efficient way to study the bias of the coin. We do, in practice, almost never consider the "catchall" hypothesis that our model is wrong. Instead, we just ask what we can learn about target system S given a certain degree of idealization and abstraction.

This observation connects to a classical problem of Bayesian inference mentioned in the introduction (see also Walker 2013; Wenmackers and Romeijn 2016; Vassend 2019). Bayesians interpret $p(H)$ standardly as the prior degree of belief that the hypothesis H is true. However, H may be part of a highly idealized statistical model of the target system that simplifies functional dependencies or neglects causally relevant factors (e.g., because the effect size

https://doi.org/10.1086/707554 Published online by Cambridge University Press

is too small to be relevant). Such idealized models occur in all fields of science and are particularly common in disciplines in which causal relationships are messy and hard to isolate, such as psychology, economics, and climate science.

In such cases, it is highly unlikely that our hypothesis H is literally true, that it sketches a faithful picture of reality. Regardless of how probable H is with respect to other hypotheses in the overarching model $\mathcal{M}$ (e.g., H′, H″), we would not be prepared to enter any bet on the truth of H. But if this is the case, how can we entertain a strictly positive degree of belief in H? In other words, why do we assign a positive degree of belief $p(\mathrm{H})$ in our Bayesian model and not probability zero?[7]

In other words, we face the trilemma of having to reject one of the following three jointly inconsistent propositions (Sprenger and Hartmann 2019, chap. 12):

1. The expression $p(\mathrm{H}) > 0$ denotes a Bayesian agent's degree of belief that H is true.
2. The hypothesis H is part of a general statistical model $\mathcal{M}$ with a partition of hypotheses $\mathcal{H} = \{\mathrm{H}, \mathrm{H}', \mathrm{H}'', \ldots\}$.
3. The model $\mathcal{M}$ is a strong idealization of reality and likely or known to be false.

The second and third propositions are commonly acknowledged facts about statistical modeling and cannot be rejected as long as we aim at a rational reconstruction of Bayesian inference in science. So we have to give up the first proposition and to rethink the interpretation of prior probabilities in Bayesian inference. When the underlying models are sufficiently idealized, they cannot denote an agent's honest degrees of belief that a hypothesis is true.

Note that the introduction of a "catchall hypothesis" $\tilde{H} = \neg(\mathrm{H} \vee \mathrm{H}' \vee \mathrm{H}'' \vee \ldots)$ into $\mathcal{M}$ (e.g., Earman 1992) does not solve the problem. Since the model $\mathcal{M}$ is highly idealized, we would have to concentrate all probability mass on $\tilde{\mathrm{H}}$. This means that we could not differentiate among H, H′, H″, and the other hypotheses in $\mathcal{M}$: their prior and posterior probability would always be zero or arbitrarily close to that value.

Extending the suppositional analysis solves the problem: if $\mathcal{M}$ describes the overarching statistical model and H is one of the hypotheses in $\mathcal{M}$, then

7. Even when the overall model is credible, it is not clear how to conduct meaningful (Bayesian) hypothesis tests. The null hypothesis H_0 usually denotes the absence of an effect, the additivity of two factors, the causal independence of two variables in a model, etc. In most cases, it is strictly speaking false: there will be some minuscule effect in the treatment, some slight deviation from additivity, some negligible causal interaction between the variables (e.g., Gallistel 2009). Yet, we would like to use it in inference and decision making, and this is difficult if its prior and posterior probability is always zero.

https://doi.org/10.1086/707554 Published online by Cambridge University Press

all probability assignments are relative to $\mathcal{M}$. The degree of belief in the truth of H that we would have if we had supposed that the target system is fully and correctly described by one of the hypotheses in $\mathcal{M}$ is expressed by $p_{\mathcal{M}}(\mathrm{H})$. In fact, Bayesian modelers distribute prior degrees of belief only over elements of $\mathcal{P}$. This perspective on prior probabilities, which flows logically from our take on conditional degrees of belief, resolves the above trilemma: we deny premise 1 that $p(\mathrm{H})$ denotes an actual degree of belief. Any such probability is, like the probabilities of evidence given a hypothesis, essentially hypothetical and should be read as $p_{\mathcal{M}}(\mathrm{H})$. Similarly, $p(\mathrm{E}|\mathrm{H})$ should always be read and understood as the model-relative conditional degree of belief $p_{\mathcal{M}}(\mathrm{E}|\mathrm{H})$.

Suppositional or hypothetical prior probabilities are, however, not rationally constrained in the same way as $p_{\mathcal{M}}(\mathrm{E}|\mathrm{H})$. In particular, supposing $\mathcal{M}$ need not define any objective chances in $\omega_{\mathcal{M},\mathcal{S}}$. Thus, unlike in the case of $p_{\mathcal{M}}(\mathrm{E}|\mathrm{H})$, the suppositional analysis does not yield a uniquely rational degree of belief for $p_{\mathcal{M}}(\mathrm{H})$ or offer any guidance "how we are supposed to understand and evaluate counterfactual probabilities," as Vassend (2019, 711) writes in a congenial contribution.

Vassend makes his observation in a critical spirit, but I would like to embrace it. Subjective Bayesian inference allows for different rational choices of the prior distribution upon supposing a model $\mathcal{M}$ and remains silent on which factors determine this choice. If $p_{\mathcal{M}}(\mathrm{H})$ were determined in the same way as $p_{\mathcal{M}}(\mathrm{E}|\mathrm{H})$, this feature of subjective Bayesian inference would get lost and we would be stuck with a specific version of objective Bayesianism. A general conceptual scheme for interpreting probabilities in Bayesian inference should not have such strong implications.

What is more, apart from resolving the above trilemma, the suppositional analysis is compatible with various strategies for interpreting and determining prior probabilities. For example, we may adopt a broadly Lewisian perspective and determine the prior probability of H as a function of the estimated similarity of ω_{H} to the actual world $\omega_{@}$ (Lewis 1973b). As an alternative, Vassend (2019) suggests that the subjective probability of H expresses the degree of belief that H is most similar to the truth among all (false) hypotheses in $\mathcal{M}$. This approach connects Bayesian inference to the verisimilitude paradigm, where the goal of scientific inference consists in gradually approaching the truth (Niiniluoto 2011; Cevolani and Tambolo 2013). Other agents will adopt an objective Bayesian approach and maximize the entropy of the probability distribution over the parameter of interest (e.g., Jaynes 1968; Williamson 2010). Finally, social conventions may dictate the choice of an "impartial" prior because any other assignment of prior probabilities would fail to convince stakeholders interested in the outcome of the statistical analysis.

These proposals are very different in their conceptual repertoire, but they all share one element: supposing a general statistical model $\mathcal{M}$ and relativizing inference to that model. For example, Vassend interprets $p(\mathrm{H})$ as the

https://doi.org/10.1086/707554 Published online by Cambridge University Press

degree of belief that H is the most truth-like among all hypotheses in $\mathcal{M}$. Objective Bayesians maximize entropy relative to a partition of hypotheses in $\mathcal{M}$. And so on. The suppositional analysis is not meant to be an alternative to these proposals but rather a general conceptual framework that accommodates various strategies for determining and interpreting prior degrees of belief.[8]

5. Implications: Model-Relative Bayesian Inference. Finally we address the question of how the suppositional analysis connects to prediction and decision making. First, we note that the above proposal naturally transfers from prior to posterior probabilities. They should be understood relative to a model $\mathcal{M}$ and be written as $p_{\mathcal{M}}(\mathrm{H}|\mathrm{E})$. Then, Bayes's theorem relates the posterior probability $p_{\mathcal{M}}(\mathrm{H}|\mathrm{E})$ to the prior probabilities $p_{\mathcal{M}}(\mathrm{H})$ and $p_{\mathcal{M}}(\mathrm{H}_i)$ and the likelihoods $p_{\mathcal{M}}(\mathrm{E}|\mathrm{H})$ and $p_{\mathcal{M}}(\mathrm{E}|\mathrm{H}_i)$:

$$p_{\mathcal{M}}(\mathrm{H}|\mathrm{E}) = \left(\sum_{\mathrm{H}_i \in \mathcal{H}} \frac{p_{\mathcal{M}}(\mathrm{E}|\mathrm{H}_i)}{p_{\mathcal{M}}(\mathrm{E}|\mathrm{H})} \cdot \frac{p_{\mathcal{M}}(\mathrm{H}_i)}{p_{\mathcal{M}}(\mathrm{H})} \right)^{-1}. \tag{1}$$

Posterior probabilities can thus only be as arbitrary as the priors are, and Bayesians of all sorts can—for nonextreme priors and sufficiently powerful experiments—use them confidently in assessing scientific hypotheses. A similar diagnosis applies to predictions of events derived from Bayesian models. According to the law of total probability,

$$p_{\mathcal{M}}(\mathrm{E}) = \sum_{\mathrm{H}_i \in \mathcal{H}} p_{\mathcal{M}}(\mathrm{E}|\mathrm{H}_i) \cdot p_{\mathcal{M}}(\mathrm{H}_i)$$

can be understood as a weighted average of the conditional probabilities of E under the competing hypothesis. For example, the hypotheses in $\mathcal{M}$ could correspond to various economic models that make different statistical predictions for specific events such as E: "On economic policy X, the Italian gross domestic product will grow for three years in a row." Similarly, some conditional degrees of belief $p(\mathrm{E}|\mathrm{H}')$ are in reality a weighted average of conditional degrees of belief because H' is a disjunction of various precise hypotheses in $\mathcal{M}$ (e.g., $\mathrm{H}' : \mu > 0$ for a parameter of interest μ).

If most hypotheses in $\mathcal{M}$ assign a high conditional probability to E (i.e., $p_{\mathcal{M}}(\mathrm{E}|\mathrm{H}_i) \approx 1$), then their weighted average $p_{\mathcal{M}}(\mathrm{E})$ will be close to 1 also. But how do these predictions affect our expectations that E will actually occur? More precisely, how should predictions of statistical models about the occurrence of future events inform real-world decisions? This question is especially

8. In particular, while Vassend and I investigate different research questions—the alignment of conditional degrees of belief with objective probability vs. the interpretation of prior probabilities—our approaches and conclusions are compatible.

https://doi.org/10.1086/707554 Published online by Cambridge University Press

important whenever there are reasons to assume that our models are highly idealized, and yet we would like to use them for making probabilistic predictions and substantiating policy making (e.g., Stainforth et al. [2007], for a case from climate science).

Indeed, for a complex and highly nonlinear target system such as the economy of a country or the earth's atmosphere, the predictions of a statistical model $\mathcal{M}$ should be taken with a grain of salt. Our confidence in its predictions depends on its grip on the causally relevant features of the target system. In other words, aligning our subjective expectation of E with $p_{\mathcal{M}}(\mathrm{E})$ depends on the adequacy of the models in $\mathcal{M}$. Correcting that value may be necessary to account for limitations of the model and neglect of particular features of the target system. This is fully analogous to non-Bayesian modeling. An engineer who calculates the trajectory of a cannonball will not blindly trust a model based on the law of gravitation and the ideal superposition of horizontal and vertical motion. In practice, factors such as air resistance and the size and form of the cannonball make a difference. The model informs our predictions, but we need to calibrate it with reality—a task that often exceeds the reach of the (physical) model and needs to be integrated with a scientist's experience and judgment. Similarly, instead of naively calibrating our actual degrees of belief with the output of a Bayesian statistical model, we should see the model as a device showing us our hypothetical degrees of belief under reasonable idealizing assumptions.

This view of Bayesian inference squares well with the famous saying from statistics that all models are wrong but some are illuminating and useful (Box 1976). Having a prior (or posterior) over μ does not commit us to any degree of belief about the "true" value of μ or to betting on some propositions about μ with specific odds. It just makes a statement about our hypothetical degrees of belief and betting odds if the coin toss were fully described by $\mathcal{M}$. This move likens inference with Bayesian statistical models to other model-based scientific inferences: one uses a (deterministic or statistical) model to derive predictions whose validity depends on the adequacy of the model itself. Moreover, both Bayesian and non-Bayesian modelers face the open question of how to transfer knowledge about the model to the target system. Whatever strategy works in the general case might also work for Bayesian models and vice versa. But answering this question in general goes beyond the scope of the suppositional analysis.

6. Conclusion. Why do probability densities of a statistical model determine the corresponding conditional degrees of belief of a Bayesian reasoner? In other words, why do we accept THE EQUALITY for observation E and statistical hypothesis H: $p(\mathrm{E}|\mathrm{H}) = \rho_{\mathrm{H}}(\mathrm{E})$?

I have argued that this question cannot be answered by recourse to ratio analysis or chance-credence coordination. In particular, the density functions ρ_{H}

https://doi.org/10.1086/707554 Published online by Cambridge University Press

do not define chances in the sense of PDI or PP. To make chance-credence coordination work, we need a different analysis of conditional degrees of belief in Bayesian inference. The suppositional analysis provided in this article answers this challenge, yields THE EQUALITY, and squares well with our intuitive handling of conditional probability. On top of this, it explains the interpretation of prior probabilities in the context of idealized models. Bayesian inference should be construed as a model-relative activity in which we reason with hypothetical, not with actual, degrees of belief. As I have argued, this impairs neither the functionality of Bayesian reasoning nor its normative pull.

The suppositional analysis has repercussions for philosophy of probability, too. If conditional degrees of belief are a central and irreducible concept in Bayesian reasoning, then we might be inclined to say the same of conditional probability. This proposal agrees with prominent axiomatizations and analyses of conditional probability as a primitive concept (e.g., Popper 1959/2002; Rényi 1970; Hájek 2003). It is up to future work to determine the most fruitful account of the relationship between conditional and unconditional probability (see also Fitelson and Hájek 2017; Gyenis, Hofer-Szabó, and Rédei 2017). Similarly, one could further explore the role of conditional probability in the various objective-chance interpretations (e.g., Frigg and Hoefer 2015; Suárez 2018).

Moreover, the suppositional analysis provides the epistemic foundations for solving central problems in confirmation theory and formal epistemology, such as the problem of old evidence (Howson 1984, 1985; Sprenger 2015) or learning information that is expressed by conditional probabilities (Sprenger and Hartmann 2019, chap. 4; Eva, Hartmann, and Rafiee Rad, forthcoming).

Finally, I recapitulate the main features of the suppositional analysis, in order to leave the reader with a coherent and unified picture of the article's results. First, conditional degrees of belief of the type $p(\mathrm{E}|\mathrm{H})$ should be interpreted in the suppositional, hypothetical way anticipated by Ramsey: we suppose that H correctly describes the target system S. The supposition is made explicit by evaluating the probability of E in a set of possible worlds $[\omega_{\mathrm{H,S}}]$, that is, the possible worlds where the behavior of S follows the probabilistic law H.

Second, chance-credence coordination principles, such as the PP or the PDI, apply even if we are agnostic about the existence of physical chance in the actual world. Instead of informing our (unconditional) degrees of belief in the actual world $\omega_{@}$, they align conditional degrees of belief with the corresponding (objective) probability densities in the $[\omega_{\mathrm{H,S}}]$ worlds. This analysis establishes THE EQUALITY and explains why statements about the probability of events in a statistical model often appear analytic. It also shows why Bayesian reasoners typically agree on Bayes factors or other measures of evidential support, thereby defending Bayesian inference against the charge of arbitrariness and lack of objectivity.

https://doi.org/10.1086/707554 Published online by Cambridge University Press

Third, all probabilities in Bayesian inference should ultimately be understood as hypothetical degrees of belief, conditional on supposing a general statistical model. This perspective explains why we are almost never willing to engage in a bet on the truth of a hypothesis in a statistical model $\mathcal{M}$. At the same time, this view does not commit us to a specific principle for determining rational prior probabilities; in fact, it sums up the common denominator of different approaches.

Fourth and last, the suppositional analysis regards Bayesian inference as a specific form of model-based reasoning in science, with features and challenges that are similar to deterministic modeling. The question of how to understand conditional and unconditional degrees of belief is not separable from the use of Bayesian models in scientific practice.

REFERENCES

Adams, Ernest W. 1975. *The Logic of Conditionals*. Dordrecht: Reidel.
Bernardo, José M., and Adrian F. M. Smith. 1994. *Bayesian Theory*. New York: Wiley.
Bovens, Luc, and Stephan Hartmann. 2003. *Bayesian Epistemology*. New York: Oxford University Press.
Box, George E. P. 1976. "Science and Statistics." *Journal of the American Statistical Association* 71:791–99.
Cevolani, Gustavo, and Luca Tambolo. 2013. "Progress as Approximation to the Truth: A Defence of the Verisimilitudinarian Approach." *Erkenntnis* 78:921–35.
Cox, David, and David Hinkley. 1974. *Theoretical Statistics*. London: Chapman & Hall.
de Finetti, Bruno. 1972. *Probability, Induction and Statistics: The Art of Guessing*. New York: Wiley.
———. 2008. *Philosophical Lectures on Probability*. Berlin: Springer.
Dubins, Lester E. 1975. "Finitely Additive Conditional Probabilities, Conglomerability and Disintegrations." *Annals of Probability* 3:89–99.
Earman, John. 1992. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.
Eva, Benjamin, Stephan Hartmann, and Soroush Rafiee Rad. Forthcoming. "Learning from Conditionals." *Mind*, https://doi.org/10.1093/mind/fzz025.
Fisher, R. A. 1956. *Statistical Methods and Scientific Inference*. New York: Hafner.
Fitelson, Branden, and Alan Hájek. 2017. "Declarations of Independence." *Synthese* 194:3979–95.
Frigg, Roman, and Carl Hoefer. 2015. "The Best Humean System for Statistical Mechanics." *Erkenntnis* 80:551–74.
Gallistel, C. R. 2009. "The Importance of Proving the Null." *Psychological Review* 116:439–53.
Gyenis, Zalán, Gábor Hofer-Szabó, and Miklós Rédei. 2017. "Conditioning Using Conditional Expectations: The Borel-Kolmogorov Paradox." *Synthese* 194:2595–630.
Hájek, Alan. 2003. "What Conditional Probability Could Not Be." *Synthese* 137:273–323.
Hoefer, Carl. 2007. "The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance." *Mind* 116:549–96.
Howson, Colin. 1984. "Bayesianism and Support by Novel Facts." *British Journal for the Philosophy of Science* 35:245–51.
———. 1985. "Some Recent Objections to the Bayesian Theory of Support." *British Journal for the Philosophy of Science* 36:305–9.
Howson, Colin, and Peter Urbach. 2006. *Scientific Reasoning: The Bayesian Approach*. 3rd ed. La Salle, IL: Open Court.
Jackson, Frank. 1991. *Conditionals*. Oxford: Oxford University Press.
Jaynes, Edwin T. 1968. "Prior Probabilities." *IEEE Transactions on Systems Science and Cybernetics* 4 (3): 227–41.

https://doi.org/10.1086/707554 Published online by Cambridge University Press

Jeffrey, Richard C. 1965. *The Logic of Decision*. Chicago: University of Chicago Press.
Jeffreys, Harold. 1961. *Theory of Probability*. 3rd ed. Oxford: Oxford University Press.
Kass, Robert E., and Adrian E. Raftery. 1995. "Bayes Factors." *Journal of the American Statistical Association* 90:773–95.
Kyburg, Henry E. 1974. *The Logical Foundations of Statistical Inference*. Dordrecht: Reidel.
Levi, Isaac. 1977. "Direct Inference." *Journal of Philosophy* 74:5–29.
———. 1980. *The Enterprise of Knowledge*. Cambridge, MA: MIT Press.
———. 1996. *For the Sake of the Argument*. Cambridge: Cambridge University Press.
Lewis, David. 1973a. "Causation." *Journal of Philosophy* 70:556–67.
———. 1973b. *Counterfactuals*. London: Blackwell.
———. 1980. "A Subjectivist's Guide to Objective Chance." In *Studies in Inductive Logic and Probability*, ed. Richard C. Jeffrey, 263–93. Berkeley: University of California Press.
McCullagh, Peter. 2002. "What Is a Statistical Model?" *Annals of Statistics* 30:1225–310.
Niiniluoto, Ilkka. 2011. "Revising Beliefs towards the Truth." *Erkenntnis* 75:165–81.
Popper, Karl R. 1959/2002. *The Logic of Scientific Discovery*. Repr. London: Routledge. Originally published in German in 1934 as *Logik der Forschung*.
Ramsey, Frank P. 1929/1990. "General Propositions and Causality." In *Philosophical Papers*, ed. D. H. Mellor, 145–63. Cambridge: Cambridge University Press.
Reichenbach, Hans. 1949. *The Theory of Probability*. Berkeley: University of California Press.
Rényi, Alfred. 1970. *Foundations of Probability*. San Francisco: Holden-Day.
Rosenthal, Jacob. 2004. *Wahrscheinlichkeiten als Tendenzen*. Paderborn: Mentis.
Savage, Leonard J. 1972. *The Foundations of Statistics*. 2nd ed. New York: Wiley.
Skyrms, Brian. 2000. *Choice and Chance*. 4th ed. Belmont, CA: Wadsworth.
Sprenger, Jan. 2010. "Statistical Inference without Frequentist Justifications." In *EPSA Epistemology and Methodology of Science: Launch of the European Philosophy of Science Association*, 289–97. Berlin: Springer.
———. 2015. "A Novel Solution of the Problem of Old Evidence." *Philosophy of Science* 82:383–401.
Sprenger, Jan, and Stephan Hartmann. 2019. *Bayesian Philosophy of Science*. Oxford: Oxford University Press.
Stainforth, D. A., M. R. Allen, E. R. Tredger, and L. A. Smith. 2007. "Confidence, Uncertainty and Decision-Support Relevance in Climate Predictions." *Philosophical Transactions of the Royal Society* A 365:2145–61.
Stalnaker, Robert. 1968. "A Theory of Conditionals." In *Studies in Logical Theory*, ed. Nicholas Rescher, 98–112. American Philosophical Quarterly Monograph Series no. 2. Oxford: Blackwell.
Suárez, Mauricio. 2018. "The Chances of Propensities." *British Journal for the Philosophy of Science* 69 (4): 1155–77.
Suppes, Patrick. 1966. "A Bayesian Approach to the Paradoxes of the Ravens." In *Aspects of Inductive Logic*, ed. Jaakko Hintikka and Patrick Suppes. Amsterdam: North-Holland.
Vassend, Olav Benjamin. 2019. "New Semantics for Bayesian Inference: The Interpretive Problem and Its Solutions." *Philosophy of Science* 86 (4): 696–718.
Walker, S. G. 2013. "Bayesian Inference with Misspecified Models." *Journal of Statistical Planning and Inference* 143:1621–33.
Wenmackers, Sylvia, and Jan-Willem Romeijn. 2016. "A New Theory about Old Evidence." *Synthese* 193:1225–50.
Williamson, Jon. 2010. *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.
Zhao, Jiaying, Anuj Shah, and Daniel Osherson. 2009. "On the Provenance of Judgments of Conditional Probability." *Cognition* 113:26–36.

https://doi.org/10.1086/707554 Published online by Cambridge University Press