Collectivist Foundations for Bayesian Statistics

Conor Mayo-Wilson and Aditya Saraf

Abstract

What (if anything) justifies the use of Bayesian statistics in science? The traditional answer is that Bayesian statistics is simply an instance of orthodox expected utility theory. Thus, Bayesian statistical methods, like principles of utility theory, are justified by norms of individual rationality. In particular, most Bayesians argue that a scientist's credences must satisfy the probability axioms if she adheres to individualistic norms of practical and epistemic rationality. We argue that, to justify Bayesian statistics as a tool for science, it is only necessary that a scientist's *public credences* (i.e., her degrees of belief qua scientist) obey the probability axioms. We call that thesis *public probabilism* and argue that, to justify it, one must appeal to norms of collective science, not just norms of individual rationality. Specifically, we argue that individualistic norms fail to justify an Archimedean condition that is necessary for credences to admit a numerical representation. However, a norm of collective science, which we call *epistemic cooperativeness*, can justify the Archimedean condition.

What (if anything) justifies the use of Bayesian statistics in science? The traditional answer is that Bayesian statistics is simply an instance of orthodox expected utility theory, and so Bayesian statistical methods, like principles of utility theory, are justified by norms of *individual rationality*. For example, one who has credences that violate the probability axioms is Dutch-Bookable, accuracy dominated, and so on. And if one is Dutch-Bookable or accuracy dominated, so the standard arguments go, then one violates some individualistic norm of practical or epistemic rationality.

In this paper, we explore an alternative defense of scientists' use of Bayesian statistical methods: norms of *collective science* may work in tandem with norms of individual rationality to justify the mathematical machinery of Bayesian statistics. We focus on *probabilism*, the thesis that experimenters' credences must obey the probability axioms.

The paper contains two major sections. In §1, we criticize (versions of) three of the most common arguments for probabilism: Dutch Book arguments, accuracy dominance arguments, and representation-theorem based arguments.[1] We argue that none of the three justifies the claim that rational credences obey an Archimedean condition. The first two merely assume credences are real numbers and so must have that property. Proponents of the third argument do try to justify the Archimedean condition but they often appeal to implausible behaviorist assumptions to do so.

In §2, we argue that, to justify the use of Bayesian statistics in science, a scientist's prior distribution cannot be interpreted as an her personal credences. Instead, a prior distribution must represent what we call the scientist's *public* credences, which are the beliefs that she adopts in virtue of her obligations as a scientist. We then argue that a norm of collective science, which we call *epistemic cooperativeness*, can justify the claim that a scientist's public credence should satisfy the Archimedean condition. Our argument provides some evidence for the claim that, together, norms of individual rationality and collective science may justify *public probabilism*, the thesis that scientists' public credences should obey the probability axioms.

Should one conclude that Robinson Crusoe may rationally violate probabilism but that scientists must heed Bayesian norms? Our arguments leave open that possibility. But our arguments are better understood as opening a different avenue for defending the claim that scientists' behavior should abide by decision-theoretic norms. To be clear, our arguments do *not* establish that the traditional, individualistic strategy for defending probabilism is hopeless; we argue only that some common individualistic arguments share a weakness. We also do not claim to establish that scientists' public credences should satisfy the probability axioms. After all, we focus on only one condition that is necessary for a probabilistic representation of credence.

The paper, therefore, should be understood as an attempted proof of concept: we aim to show that collective norms for scientific inquiry might better justify what are often taken as norms for rational, individual behavior. Equally importantly, although we focus on probabilism, our argumentative strategy may be used to justify other parts of Bayesian statistical practice that, we believe, currently lack cogent philosophical foundations. For example, some Bayesian statisticians argue that one ought to "choose" noninformative priors when analyzing data; others argue that scientists ought to choose priors that have good frequentist properties.[2] From the standpoint of traditional decision theory, analyzing data using a probability distribution

[1] For a summary of Dutch book arguments, see [Vineberg, 2016] and references therein. For accuracy dominance arguments, see [Joyce, 1998] and [Pettigrew, 2016]. For use of representation theorems, see [Savage, 1972] and [Krantz et al., 2006].

[2] See Excursion/Chapter 6 of [Mayo, 2018] for a discussion of various ways in which contemporary Bayesians choose priors.

that deviates from one's personal credences is typically irrational. Just as it would be irrational for Alison to choose an action just because it maximizes Bill's subjective expected utility, a Bayesian statistician who chooses a prior in one of the two ways just specified may advocate statistical decisions that would fail to minimize expected loss by her own lights. This paper suggests a novel way to rationalize these two methods for "choosing" a prior: norms of science might require a statistician to adopt public credences that differ from her private ones.

1 Individualistic Foundations of Bayesianism

The fundamental thesis of Bayesianism is *probabilism*, which asserts that one's credences over an algebra of events $\mathcal{A}$ should be representable by a probability measure $P : \mathcal{A} \to [0, 1]$. What does it mean for credences to be representable by a probability measure? Philosophers and social scientists typically assume an agent's credences are represented by a binary relation $\succeq$ on an algebra of events, where $A \succeq B$ represents the claim that the agent believes A to be at least as likely as B.[3] To say that an agent's credences are "probabilistic" (or representable by a probability measure), therefore, is to say there is at least one probability measure P such that $A \succeq B$ if and only if $P(A) \geq P(B)$.[4]

Several canonical arguments for probabilism simply assume that credences are numerically representable, i.e., that there is a *real-valued* function $r : \mathcal{A} \to \mathbb{R}$ such that $r(A) \geq r(B)$ if and only if $A \succeq B$. Dutch Book arguments, for example, do so by identifying credences with fair prices (see next section), and accuracy-dominance arguments often contain numerical representability as a modeling assumption. Proponents of those arguments,

[3]Historically, binary representations are the most common way of representing comparative credence. However, it is well-known that some aspects of human perception are best represented by higher arity relations. see [Krantz et al., 2006, p. 139]'s discussion of difference structures. Further, choice behavior is often represented by *choice functions* [Seidenfeld et al., 2009, Sen, 1971], which take sets of arbitrary size and output the options that are "choice-worthy." So it is not beyond the realm of possibility that credence should similarly be modeled by relations of higher arity or some mathematical apparatus like choice functions. We focus on binary relations for simplicity.

[4]What we call representability is often called "full agreement"; see [Fishburn, 1986]. Our argument in §1.2 requires only the conjunction of what Fishburn calls "almost" and "partial agreement", and that conjunction is not equivalent to full agreement except in the presence of the additional assumption that $\preceq$ is total. This distinction matters because some decision theorists aim only to defend the claim that rational credence agrees with a probability measure in one of these weaker senses.

therefore, focus on defending the claim that numerically-representable credences obey the remaining probability axioms (i.e, additivity and normality). In this paper, we take a step back and question what (if anything) justifies the assumption that credences are representable by real numbers. In §1.1, we first argue that some descriptive psychological theories of credence (specifically, behavioristic ones) fail to justify that assumption, and in §1.2, we argue that normative theories of (individualistic) rationality likewise falter.

1.1 What is credence?

There are roughly two types of theories of credence: behaviorist and mentalist.[5] Famously, behaviorists argue that, to explain and predict human behavior, it is unscientific to postulate the existence of "internal" mental states. Thus, some behaviorists *identify* psychological states, including desires and beliefs, with observable behavioral dispositions.

For example, a behaviorist might identify an agent's credence $c(A)$ in A with her fair price $Pr(A)$, i.e., the price the agent is willing both to buy and sell a gamble that is worth \$1 if A occurs and is worth \$0 otherwise. Other behaviorists claim that an agent's credence in B is *stronger* than her credence in A (written $A \prec B$) if, when forced to choose between (i) receiving a prize if A occurs or (ii) receiving the same prize if B occurs, the agent would choose the latter option. According to this second account, credence is a disposition to make certain choices when decisions are *forced*.

Most philosophers and social scientists reject these behaviorist accounts of credence.[6] If credence is nothing more than the disposition to gamble in particular ways in laboratory settings, then it is of little philosophical interest. When philosophers and social scientists discuss credence, their working hypothesis is that credence is an underlying mental state that is causally related to many other mental states (e.g., desires and regrets) and behaviors.[7] Following [Okasha, 2016], we call such views of credence *mentalist*.

According to mentalism, fair prices are one crude way of measuring credence, but they are not the only way. Further, fair prices can be wildly misleading: when forced to offer prices in a laboratory, a subject's behavior may not reflect any underlying attitude that can explain or predict her choices in other scenarios. Similar criticisms apply to the behaviorist account that identifies credence with forced choices.

[5] For an enlightening discussion of behaviorist and mentalist views, see [Okasha, 2016].

[6] See [Eriksson and Hájek, 2007] for trenchant criticisms of behaviorist views.

[7] For example, see [Walley, 1991, §1.3]'s discussion of "theoretical" interpretations of probability.

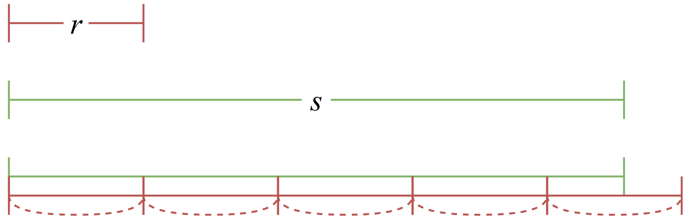

$5r > s$

Figure 1: An example of the Archimedean property for real numbers.

We will not defend mentalism, but we do reject the two behaviorist theories above for the reasons sketched. We criticize those accounts to anticipate an objection to our main argument. According to some behaviorist theories, credence, *whether rational or not*, is representable by a numerical (real-valued) function, and so it follows that rational credence obeys an Archimedean condition, pace our thesis. For example, according to the first account, credences are fair prices, which are, by definition, numerical.

There are, of course, many ways that one might try to identify credence with behavior. We cannot prove that no such attempt will be successful. Rather, we claim only that if one identifies credence with behavior in a way that make it easy to justify the claim that credences are numerical (and hence, satisfy an Archimedean condition), then one typically faces the difficult task of defending the claim that credence, so defined, plays any important predictive or explanatory purpose in describing *other* behaviors.

In the rest of the paper, we simply assume that a person's credence might not be real-valued. The central question of the next section is whether *rational* credence must be numerical. We argue not.

1.2 Archimedean Conditions

Real numbers satisfy an Archimedean property: for any two numbers $r, s > 0$, there is some natural number n such that $n \cdot r > s$. See Figure 1 for a visual example. Thus, if credences are numerically representable, then they must obey an Archimedean-like condition. In this section, we formulate a (simple) comparative Archimedean condition — one that can be expressed solely in terms of $\preceq$ and the set-theoretic structure among events in an algebra. Then, we provide an example of comparative credences that seem rational but violate the condition. Thus, our Archimedean condition seems not to be justified by norms of individual rationality.

We first sketch our argument. Many philosophers (e.g., Nicolas of Cusa) have believed that space is infinite and that the number of planets is likewise infinite. Suppose, for the sake of argument, that Nicolas of Cusa contemplated the following two hypotheses: H_1, exactly one planet, Earth, contains water in liquid or solid form on its surface, and H_2, at least two planets contain water. For each positive number n, let W_n denote the proposition 'The n^{th} planet from Earth is the *closest* planet that contains water." So H_2 may be thought of as the infinite disjunction of all W_ns.

We claim it is rationally permissible for Cusa to believe that, for sufficiently large numbers n, the proposition W_n has positive probability. Why? There is some chance that very distant planets contain water, and if so, there must be a closest such planet. Further, for sufficiently large numbers n and k, Cusa may permissibly believe that W_n and W_k as equally likely, since Cusa knows nothing that would distinguish very far away planets from one another. Yet if Cusa holds such beliefs, we claim that his comparative credences are not real-valued. To see why, consider the following condition.

Qualitative weak Archimedean condition for credence (QWAC): Every bounded and disjoint sequence of events is finite in length, where a sequence $A_1, A_2, \ldots$ is *bounded and disjoint* if it has three properties:

1. $A_i \cap A_j = \emptyset$ if $i \neq j$,
2. $A_i \succeq A_1$ for all i, and
3. $A_1 \succ \emptyset$.

Why is QWAC to the Archimedean condition for real numbers? Think of the events $A_1, A_2, \ldots$ as line segments, just as real numbers were depicted by line segments above. Condition 1 says the line segments do not overlap; condition 2 says all of the segments are at least as long the first one, and condition 3 says the first segment has positive length. So if there were an infinite sequence of bounded and disjoint events $A_1, A_2, \ldots$, then its union $B = \bigcup_{n \in \mathbb{N}} A_n$ would be representable by an infinite line and would be "infinitely more probable" than A_1.

If an an agent's comparative credence relation $\preceq$ is representable by a finitely additive probability measure P, then $\preceq$ satisfies QWAC.[8] Thus, if it is rationally permissible for one's credences $\preceq$ to violate QWAC, then it is rationally permissible to have credences that violate the probability axioms. We now defend the antecedent of that conditional.

[8]Proof: Otherwise, there is an infinite bounded and disjoint sequence $A_1, A_2, \ldots$, and

Suppose Nicolas of Cusa believes there are infinitely many planets other than Earth, and imagine he enumerates those planets 1, 2, etc., in terms of their distance from Earth.[9] For each $n \geq 1$, let W_n be the proposition that "Planet n is the *first* non-Earth planet in Cusa's enumeration that contains water." The events W_n and W_m are disjoint if $m \neq n$ because two different planets cannot both be the first in Cusa's enumeration to contain water.

Intuitively, it is rationally permissible to believe that some far away planet other than Earth contains water. Thus, we claim it is rationally permissible for Cusa to regard W_{n_0} as strictly more probable than the contradictory event $\emptyset$ for some sufficiently large number n_0. For example, Cusa might reason as follows: "I have no knowledge whatsoever of the 137^{th} planet from Earth, and it is possible that said planet contains water on its surface." Let $A_1 = W_{n_0}$, and in general, $A_n = W_{n+n_0}$. Because W_m and W_n are disjoint when $m \neq n$, so are A_m and A_n. Hence, $\langle A_n \rangle_{n \in \mathbb{N}}$ satisfies the first condition of a bounded and disjoint sequence. Because Cusa regards $A_1 = W_{n_0}$ as strictly more likely than $\emptyset$, the third condition is also satisfied.

Finally, because n_0 is very large, it is rationally permissible, we think, to regard the following two claims as equally likely: "Planet n_0 is the first planet in Cusa's enumeration to contain water" and "Planet $n + n_0$ is the first planet in Cusa's enumeration to contain water", where n is any natural number. After all, what Cusa knows about very distant planets does not distinguish planet n_0 from planet $n + n_0$. Thus, we believe it is rationally permissible for Cusa to regard A_n to be as likely as A_1 for all n. Hence, Cusa's credences in the A_n's satisfy the three conditions of a bounded and disjoint sequence, and because Cusa believes there are infinitely many planets, the sequence $\langle A_n \rangle_{n \in \mathbb{N}}$ is infinite. Thus, Cusa's credences violate the Archimedean condition above, and we have argued there is nothing rationally impermissible about those credences.

One might object that it is irrational to regard $A_{10^{100}}$ as equiprobable as A_1. Why? Imagine an explorer sets forth from Earth searching for planets

---

thus for all n:

$$\begin{aligned} P(\bigcup_{k \leq n} A_k) &= \sum_{k \leq n} P(A_k) \text{ as } A_i \cap A_j = \emptyset \text{ if } i \neq j \\ &\geq n \cdot P(A_1) \text{ as } A_1 \preceq A_i \text{ for all } i \end{aligned}$$

Since $\emptyset \prec A_1$, we know $P(A_1) = r > 0$. Hence, by the Archimedean property of real numbers, there is some sufficiently large n such that $P(\bigcup_{k \leq n} A_k) \geq n \cdot r > 1$, which is a contradiction.

[9]Cusa does not consider the possibility that the number of planets is uncountable. Nor should he for physical reasons.

with water. If A_{10} were as probable as $A_{10^{100}}$, then the explorer would be just as likely to encounter ten consecutive waterless planets as she would be to encounter a googolplex of waterless planets.[10] Finding that consequence unpalatable, one might argue that there must be *some* sufficiently large number m_1 such that Cusa regards A_{m_1} as less probable than A_1. And if one repeats the same reasoning, one would conclude there must be some sufficiently large number $m_2 > m_1$ such that Cusa regards A_{m_2} as less probable than A_{m_1}. And so on. If the sequence $A_1, A_2, \ldots$ decreases in probability in this way, then QWAC is satisfied.

The problem with this objection is that, for a credal ordering to be representable as a probability function, it is not sufficient for it to obey QWAC (even in conjunction with other standard "axioms" for credence). To see why, note that although the sequence $\frac{1}{2}, \frac{1}{3}, \frac{1}{4} \ldots$ is decreasing, its sum $\sum_{n \geq 2} \frac{1}{n} = \infty$ diverges. In general, even if a real-valued function $P(A_n)$ decreases strictly with each n, if the terms do not decrease *fast enough*, then P is not a probability measure. Thus, any Archimedean condition that is sufficient for probabilism must require not only that Cusa's credences in A_n decrease, but also that they decrease *at a specific rate*. We think it is implausible that such a precise constraint is a requirement of rationality. To keep our exposition brief, we will not state an Archimedean condition that is sufficient for real-valued representation; see [Fishburn, 1986, p. 342] for several versions.

2 Collectivist Foundations for Bayesianism

2.1 Public Probabilism

In the previous section, we argued that several arguments for probabilism — which appeal to behaviorist assumptions or to norms of individual rationality — share a common weakness: they do not justify QWAC. We now argue that, luckily, to justify the use of Bayesian statistics in science, one need not assume that a scientist's *private* credences satisfy the probability axioms.

Our view is not entirely novel. Many statisticians already reject the view that the prior probability distributions appearing in scientific journals should be interpreted as the author's credences. For example, [Gelman and Shalizi, 2013, p. 19], two prominent statisticians, write, "... the prior distribution $p(\theta)$ is one of the assumptions of the model and does not need to represent

[10] This objection is inspired by an objection raised by an anonymous referee, whose comments significantly improved this paper.

the statistician's personal degree of belief in alternative parameter values." They continue:

> We do not have to worry about making our prior distributions match our subjective beliefs still less about our model containing all possible truths. Instead we make some assumptions, state them clearly, see what they imply, and check the implications. This applies just [as] much to the prior distribution as it does to the parts of the model showing up in the likelihood function [Gelman and Shalizi, 2013, p. 20].

We will not reconstruct Gelman and Shalizi's view in detail, but the quotation raises an interesting possibility: Bayesian statistical methods might be justified without discussing rational belief at all!

Unlike Gelman and Shalizi, we do think of prior probabilities as representing belief *in some sense*, but not as representing the experimenter's *actual* credences, nor even as idealizations of those credences. Instead, prior probabilities are, we claim, best interpreted as a model of a scientist's *public* (or "professional") credences.

We think of modeling scientists' beliefs in much the same way one might model the beliefs of a juror in a criminal trial. Famously, jurors are required to ignore some types of evidence (e.g., hearsay), and they are obliged to consider other types (e.g., exhibits introduced during trial). A model of a juror's decision-making — whether descriptive or normative — could be wildly inaccurate if it considered only the private beliefs of the juror qua citizen; what matters are the beliefs of the juror qua juror. We refer to these as the juror's *public* beliefs. The juror's public beliefs are, we think, rightfully called "beliefs" because they explain the juror's courtroom behavior in much the way her private beliefs explain her behavior in her private life.[11]

Analogously, participation in a scientific community requires one to adopt certain beliefs and to modify those beliefs in a particular way, regardless of one's private opinions. A model of scientific decision-making — whether normative or descriptive — might likewise be inaccurate if it confuses the beliefs of a scientist qua scientist with her beliefs qua private citizen.

[11] Nothing in our argument hinges on us calling the juror's attitudes "beliefs." Some philosophers might prefer to say the juror "accepts" certain propositions [Cohen, 1992, Levi, 1967, Maher, 1993, Van Fraassen, 1980]. For those who prefer the language of acceptance, our thesis is that scientists have some professional, graded acceptance-like attitude that must obey the probability axioms. [Fleisher, 2018] defends a similar thesis, namely, that the hypotheses a scientists *endorses* should be selected to maximize expected (epistemic) utility. Comparing our view to Fleisher's is beyond the scope of this paper.

For example, consider the paleontologist Marcus Ross, who identifies as a creationist.[12] Ross professes to believe that the universe is only 10,000 years old despite having written a dissertation on a marine animal that is widely accepted to be 65 million year old. When asked whether the reasoning in his dissertation was sound, Ross responded, "Within the context of old age and evolutionary theory, yes. But if the parameter is different, portions of it could be completely in error." Here, Ross distinguishes between his private beliefs and what he advocates on the pages of scientific journals. Ross' personal religious convictions, therefore, might play a relatively unimportant role in explaining or predicting his professional behavior.

Similarly, any attempt to *evaluate* Ross' scientific work should not confuse his scientific reasoning with his religious thinking. Ross' publications, grant proposals, and teaching — his behavior qua scientist — might be scientifically rigorous and satisfy all the demands of normative decision theory. Of course, Ross' behavior considered as a whole may be deeply irrational: his professional choices seem incoherent given his private beliefs. Yet that irrationality is irrelevant to the coherence of his professional behavior.

As Ross' work illustrates, a scientist's "public" beliefs may diverge from her private beliefs because she is obliged to act is if she endorses the core background assumptions of her field. Paleontologists, for instance, must act as they believe that the Earth is more than 10,000 years old.

Because a scientist's private credences may differ from her public ones (i.e., her credences qua scientist), probabilism — if understood as a thesis about a scientist's private credences — is neither necessary nor sufficient to justify the use of Bayesian statistics in science. Instead, the use of Bayesian statistics requires the truth of what we call *public* probabilism, which asserts that a scientist's public credences should obey the probability axioms, at least on a restricted set of propositions in her domain of expertise.

Although we have thusfar focused on a scientist's public beliefs at a fixed time, we emphasize that Bayesianism and other epistemological theories typically contain norms (e.g., conditionalization) about how one should *modify* one's beliefs in response to evidence. To justify the use of Bayesian methods in science, therefore, it is again not sufficient to argue that a scientist's private credences are updated by conditionalization: one must argue that a scientist's public credences are changed in that way.

Luckily, there are good reasons to believe diachronic epistemological norms might likewise be justified by norms of collective science. In general, when we participate in different social institutions (e.g., as jurors, journal

[12] Our description of Ross's beliefs is based on [Rosin, 2007]'s report.

referees, etc.), there are two reason we may be required to respond to evidence in ways we would not as private individuals. First, the evidence that we are required to consider (or not) in virtue of the institution's goals may differ from the evidence that is available to us as individuals. Second and most importantly for our purposes, procedural constraints require us to *evaluate* evidence in particular ways (e.g., rubrics in hiring committees). We now elaborate one such procedural constraint, which we call "epistemic cooperativeness", and we explain how it can be used to justify QWAC.

2.2 Epistemic Cooperativeness

In the next section, we argue that a norm of scientific inquiry, which we call *epistemic cooperativeness* (or just "cooperativeness" for short), requires that scientists' public credences obey QWAC. In this section, we define "cooperativeness" and argue it is a norm of contemporary science.[13]

Epistemic cooperativeness, in our sense, is a relation between a researcher, a scientific question, and a community, which may include both experts and laypersons. Roughly, we say a researcher is epistemically cooperative (or just "cooperative", for short) if she is *publicly* open-minded towards all "live" hypotheses about the question at hand. For example, Galileo was cooperative towards his peers and the general public with respect to the question of whether a heliocentric or geocentric model best explained everyday phenomena and astronomical data at his time. Whether Galileo privately believed that geocentrism was unfathomably stupid is irrelevant to whether he was cooperative in our sense. Galileo's published writings and his correspondences with Church officials contain a serious engagement with geocentrism.

We emphasize three features of our definition of cooperativeness. First, what counts as "live" can change; astronomers no longer need to test geocentric models. Second, a scientist may be cooperative but extremely critical of her peers' work. Third, cooperativeness does not require "impartiality" on "non-partisanship." A cooperative scientist may seek to discredit others and to confirm her own pet hypotheses. In fact, an epistemically cooperative scientist may be motivated entirely by fame, monetary prizes, and personal grudges against her rivals.[14]

[13] Although sociologists of science have studied norms of science extensively (e.g., see "The Normative Structure of Science" in [Merton, 1973]), our notion of "epistemically cooperative" cuts across several different purported scientific norms (e.g., Merton's norms of "disinterestedness", "universalism", and "communism"), and so we give it a new name.

[14] We take no stance on whether such behavior or motivations violate other scientific,

What precisely is required to be cooperative in our sense? Here is a very weark condition that we think is necessary (but insufficient).

Weak, Diachronic Credal Cooperativeness: (DCC): If H_1 and H_2 are live hypotheses and a scientist receives an indefinite amount of evidence that (a) is *for* H_1 and (b) *favors* H_1 over H_2, then she must eventually *not* regard H_2 as infinitely more likely than H_1.[15]

The reader might wonder why what we call "cooperativeness" is a norm that arises only for scientists engaged in collective inquiry. Why would an inquisitive person deserted on a remote island not be bound by the same norms? Why not call the virtues we have described "open-mindedness" or "vigilance against error", which are phrases that seem to describe virtues that a researcher can exhibit in isolation?

We do not deny the existence of purely individualistic epistemic norms. We claim only that those norms, if they exist, are weaker than what many philosophers have imagined. To see why, consider the question, "*Which* hypotheses should a scientist take seriously when evaluating data?" Clearly, a scientist is not obliged to consider *all* hypotheses that explain her data for there are innumerable such hypotheses that have yet to be articulated.

We claim that the set of hypotheses that a scientist is obliged to consider depends upon, among other things, the history of her field of study and the current work of her peers. Why? A scientist would be considered negligent if, each time she acquired novel data or designed a new experiment, she ignored all past research and considered only several new hypotheses that she personally devised. Practically speaking, the norm to consider rival hypotheses is often enforced through peer review, where referees for grant-giving institutions and journals evaluate whether the proposer or author has discussed and cited "relevant" or "appropriate" literature."[16] Similar prac-

moral, or social norms. However, Kitcher [1990, 1995] and Strevens [2003] argue that "epistemically impure" motives might improve science. For criticisms of those arguments, see [Zollman, 2018] [Heesen, 2018] and [O'Connor and Bruner, 2019].

[15] Here, we follow likelihoodists like [Edwards, 1984] and [Royall, 1997], who use the word "favors" to denote a ternary relation between some datum and two hypotheses. Note a datum may "favor" one hypothesis over another but be poor evidence for both. This is why the first condition of DCC requires that E is evidence *for* H_1. For the moment, we leave the notions of "evidence for" and "favoring" undefined, but we discuss some theories below. Equipped with such theories, we can say a sequence of propositions $\langle E_i \rangle_{i \in I}$ is "indefinite" evidence for H_1 if (1) E_i is evidence for H_1 for all $i \in I$, (2) I is infinite, and (3) the E_is are mutually conditionally independent given every $\theta_j \in H_j$ for $j = 1, 2$, where θ_j is a disjunct that makes H_j true. Define "indefinite favoring" similarly.

[16] For example, referees for *Nature* are asked, "does this manuscript reference previous

tices show many scientists' obligations depend upon their institutional roles (e.g., as employees of public universities, members of academic societies, journal referees, etc.) and upon the interests of society at large.

One might object that, although cooperativeness is a collective norm, DCC is not. The objector might grant that the set of hypotheses that a scientist must consider is determined (in part) by professional and social obligations, but that once that set is determined, DCC amounts to the duty to be responsive to evidence. Such a demand is a paradigmatic norm of individual (epistemic) rationality.

We disagree. As we show in the next section, DCC entails that a scientist's credences should obey QWAC. Thus, if a rational agent's credences may violate QWAC, then DCC is not a requirement of individual rationality. For example, we have argued that Nicholas of Cusa is individually rational, but we will argue he is not weakly cooperative, if the hypothesis that "Earth is the only planet with water" is live. The objector, we think, mistakes a vague, individualistic epistemic norm (to consider evidence) with a rather precise, formal consequence of cooperativeness.

At the very least, we think the normative force of DCC is *strengthened* by a scientist's obligations towards others. Imagine an uncooperative scientist S who *publicly* believes a hypothesis H and might continue to believe H (publicly) in the face of indefinite evidence to the contrary. In such a case, S's colleagues would be unlikely to seek out her paper defending H, for her colleagues would know that S might conclude that H despite large amounts of disconfirming evidence; after all, S's dogmatic stance towards H is *public*.

If S's data were publicly accessible, then expert readers could evaluate the data themselves. But data are not always publicly accessible, and reanalyzing another researcher's work is often costly and time-consuming. Equally importantly, some scientists may need to rely on S's results but lack the technical knowledge to re-analyze S's data. Finally, even if S honestly discloses her data, her colleagues might worry that she ignored or failed to report relevant evidence because, by supposition, some pieces of evidence will not sway S at all. In short, if a scientist is uncooperative in our sense, then her research is likely to be ignored by her colleagues, and if it is not ignored, it might not be trusted. Thus, even if DCC is partially supported by norms of individualistic rationality, its force is strengthened by considering a scientist's obligations to her peers.

Is cooperativeness so described actually a norm of science? As we have

literature appropriately? If not, what references should be included or excluded?" [Board, 2020].

discussed, peer review (of articles and grant proposals) suggests that scientists are required to evaluate how well their evidence bears on a variety of live hypotheses. Contemporary calls for pre-registered trials — with the requirement to include a detailed data-analysis plans — provide further evidence that scientists are discouraged from using methods that preclude the possibility of finding support for their rivals' hypotheses [Nosek et al., 2018].

2.3 Non-Archimedean credences violate epistemic cooperativeness

We now argue that if a scientist initially regards one hypothesis H_2 as "infinitely more probable" than another H_1 — in other words, her prior credences over hypotheses violate the Archimedean condition — then she might acquire an indefinite amount of data favoring H_1 and nonetheless continue to believe that H_2 is infinitely more plausible than H_1.[17] Thus, her beliefs violate DCC. To explain our argument, we rely on the example involving Nicolas of Cusa, but by appropriate substitutions, our formal argument generalizes to any case in which an experimenter has non-Archimedean credences in the sense described in §1.2

As before, suppose Cusa regards the hypothesis H_2 "At least two planets contain water" as infinitely more probable than H_1, "Only Earth contains water." More precisely, assume H_2 can be understood as the disjunction $H_2 = \bigcup_{n \in \mathbb{N}} A_n$ of infinitely many hypotheses, where A_n asserts that "Planet n is the *first* planet in the enumeration that contains water." We formalized the idea that Cusa regards H_2 as "infinitely more probable" than H_1 by assuming that the propositions $A_1, A_2, \ldots$ are at least as probable as H_1. Since the A_n's form an infinite, disjoint sequence, this means Cusa's credences violate the comparative Archimedean condition we discussed.[18]

Now imagine that, on successive discrete stages of inquiry — call those stages 1, 2, and so on — Cusa learns the proposition E_n, "Planet n does *not* contain water." After n many stages, Cusa will know $\bigcap_{k \leq n} E_k$, which is the claim "None of the non-Earth planets numbered less than n contain water." Notice $\bigcap_{k \leq n} E_k = \bigcap_{k \leq n} \neg A_k$ because both $\bigcap_{k \leq n} E_k$ and $\bigcap_{k \leq n} \neg A_k$

[17] Some philosophers have argued that being appropriately "open-minded" *requires* that one's credences be representable by hyperreals, which violates an Archimedean condition. See [Easwaran, 2014] for a summary and criticism of those arguments.

[18] Recall, in the previous section we defined A_k to be the proposition "The $(n_0 + k)^{th}$ planet from Earth contains water", where n_0 is a number that is large enough to guarantee that Cusa has no knowledge distinguishing planets numbered n_0 or greater. To simplify notation, we let $n_0 = 0$ in the remainder of the paper.

assert that none of the first n many planets contain water. So on stage n, Cusa believes

$$F_n = H_1 \cup \bigcup_{k \geq n+1} H_2 \cap A_k$$

which is the proposition "Either only Earth contains water or the first non-Earth planet that contains water must be numbered $n+1$ or larger."

Thus, on stage n, Cusa's credences in H_1 and H_2 on stage n should be represented by the expressions $H_1|F_n$ and $H_2|F_n$ respectively. We now argue that Cusa still regards H_2 as "infinitely more probable" than H_1 *at every stage* n, thus violating DCC. More precisely, we prove that, *conditional on* F_n, the sequence $H_1, A_{n+1}, A_{n+2}, \ldots$ is infinite, bounded and disjoint under the following assumption about Cusa's credal relation $\preceq$.

Assumption: For any piece of evidence E and any hypotheses H and H', if (a) $H \cap E \preceq H' \cap E$ and (b) $H|E$ and $H'|E$ are well-defined, then $H|E \preceq H'|E$. If the inequality in (a) is strict, then $H|E \prec H'|E$.

The assumption is somewhat controversial. It asserts that, at *every* stage of inquiry, Cusa's posterior credences are constrained by his initial/prior credences.[19] That is true if Cusa's prior credences are representable by a numerical probability function P and, for any hypothesis H and datum E, his posterior credence $P(H|E)$ in H after learning E is defined by the ratio formula $P(H|E) = \frac{P(H \cap E)}{P(E)}$. In other words, the assumption is true if Cusa is a Bayesian. However, we think it is plausible even if Cusa's credences are not probabilistic and if the axioms governing comparative/qualitative probability permit conditioning on events with zero probability.

We now prove that, at every stage of inquiry, Cusa still regards H_2 as infinitely more probable than H_1.

Claim: $A_{k+n}|F_n \succeq H_1|F_n$ for all n and all $k \geq 1$. Further, $H_1|F_n \succ \emptyset|F_n$, and so the sequence $H_1, A_{n+1}, A_{n+2}, \ldots$ is a bounded disjoint sequence conditional on F_n.

Proof: To prove the first claim, note that by assumption, $A_m \succeq H_1$ for all m. Thus, $A_{k+n} \succeq H_1$ for all n and all $k \geq 1$. Since $H_1 \cap F_n = H_1$, it follows that $A_{k+n} \succeq H_1 = H_1 \cap F_n$ for all n and all $k \geq 1$. Similarly, $A_{k+n} \cap F_n = A_{k+n}$ as $A_{k+n} \subseteq F_n$. Thus, $A_{k+n} \cap F_n \succeq H_1 \cap F_n$. By the assumption, it follows that $A_{k+n}|F_n \succeq H_1|F_n$ for all n and all $k \geq 1$.

[19] In other words, it entails what [Levi, 1983, p. 82] calls "confirmational tenacity." Levi rejects the principle as a type of dogmatism.

To show the second claim, note $\emptyset \cap F_n = \emptyset \prec H_1 = H_1 \cap F_n$ by assumption that $H_1, A_1, \ldots$ is a bounded disjoint sequence (specifically, the third condition). So by the assumption above, $\emptyset | F_n \prec H_1 | F_n$ as desired. $\square$

To sum up, in §1.2, we argued that Nicolas of Cusa can be rational while holding the private belief H_2, that "At least two planets contain water", no many how many waterless planets he observes. But if his peers regard hypothesis H_1 "Only Earth contains water" as a legitimate competitor to H_2, then Cusa violates the norm of epistemic cooperativeness if he collects only data that he knows will never lead him to take his peers' view seriously in his professional life.

One might object that, in the example given, Cusa never acquires evidence that favors H_1 over H_2. After all, each new observation of a waterless planet is still compatible with there being at least one distant planet containing water. And if no observation E_n favors H_1 over H_2, then Cusa does not violate DCC. Further, one might object that no datum could really count as evidence *for* H_1 because, if there really are infinitely many planets, then H_1 is unverifiable.

These objections are not without merit. In fact, several theories of favoring and evidence vindicate the objections. For example, according to one version of likelihoodism [Edwards, 1984, Royall, 1997], a datum E favors one *simple/atomic* hypothesis θ_1 over another θ_2 when $P_{\theta_1}(E) > P_{\theta_2}(E)$, and E favors a *composite/disjunctive* hypothesis H_1 over H_2 if and only if E favors θ_1 over θ_2 for *every* disjunct θ_1 of H_1 and every disjunct θ_2 of H_2. In our example, a qualitative version of such a likelihoodist theory would say that E_n favors the one-planet hypothesis H_1 over the alternative H_2 if and only if $E_n | H_1 \succ E_n | A_k$ *for every* k. But since both H_1 and A_{n+1} *entail* E_n (as both entail the n^{th} planet does not contain water), one might reject the claim that E_n is more probable under H_1 than under A_{n+1}.

Our response is that the running example shows the limitations of such a theory of favoring. While each observation of a waterless planet is compatible with *some way* that H_2 might be true, each also *refutes* infinitely many other ways H_2 could have been true (specifically, each E_n refutes A_n, and there are infinitely many ways that the n^{th} planet might be the first non-Earth planet with water). Further, each E_n is *entailed* by H_1. For these reasons, we suspect that many readers will, like us, believe the conjunction $\bigcap_{k \leq n} E_k$ favors H_1 over H_2 in some meaningful sense.

Small modifications to the likelihoodist thesis above support that intuition. For instance, one might argue that E favors a composite/disjunctive hypothesis H_1 over H_2 if and only if (1) E favors *at least one* disjunct θ_1

of H_1 over at least one disjunct θ_2 of H_2 and (2) no pair of disjuncts of the two hypotheses yield a favoring relation in the reverse direction. By that criterion, each datum E_n favors H_1 over H_2 in our main example.

Minor variants of Bayesian confirmation theory also entail that (a) at least one datum E_n is evidence for H_1 and (b) no E_k is evidence against H_1. For example, a confirmation theorist might argue that E confirms H if (1) $P(H|E) > P(H)$ for at least one rational prior P that assigns H positive probability and (2) there is no rational prior P that assigns H positive probability such that $P(H|E) < P(H)$.[20] Because we have argued that rational credence need not be probabilistic, clearly we do not endorse such a theory of evidence. But conditions (1) and (2) have plausible, purely qualitative analogs,[21] and with plausible assumptions about rational credal orderings, those qualitative analogs likewise entail that least one datum E_n is evidence for H_1 and no E_n is evidence against H.

The main limitation of our argument, as noted in §1.2, is that QWAC is not sufficient for a credal relation to admit a probabilistic representation, even in the presence of other plausible constraints on a rational credal relation (e.g., transitivity, additivity, and so on). Despite this limitation, we think the above argument shows how norms of collective science might be used to justify, at least partially, comparative axioms that are necessary for an experimenter's public credences to obey the probability axioms. Importantly, we do not claim that norms of collective science are sufficient, by themselves, to justify public probabilism; norms of individual rationality might still have some role to play.

2.4 Comparison to other norms

Our weak, diachronic, credal cooperaticness norm (DCC) may appear similar to other "open-mindedness" norms defended by formal epistemologists and philosophers of science. For example, some argue that rational credence should be represented by a *regular* probability measure, i.e., a function P such that $P(A) > 0$ whenever A is non-empty. So as not to presuppose probabilism, we can rephrase regularity as follows: if $\preceq$ is a rational credal

[20] If rational priors are countably additive, then we can replace (i) and (ii) by the requirement that $P(H|E) > P(H)$ for *every* rational prior P. Why? If P is countably additive, then since $P(H_2) > 0$, it follows that $P(A_n) > 0$ for some n. Thus, if Cusa learns E_n, he rules out a disjunct of H_2 that had positive prior probability. As $P(H_2) < 1$, it follows that $P(H_2|E_n) < P(H_2)$. Since $P(H_1|E_n) = 1 - P(H_2|E_n)$ and $P(H_1) = 1 - P(H_2)$, the result follows.

[21] For example, the qualitative analog of condition 1 is that $H|E \succ H$ for every rational credal ordering $\preceq$ such that $H \succ \emptyset$.

ordering, then $A \succ \emptyset$ for any $A \neq \emptyset$. Of course, our justification for DCC differs from existing arguments for regularity and similar "open-mindedness" requirements, as the latter are intended to be requirements of individual rationality. But there are two further differences between DCC and regularity.

First, DCC is a diachronic condition that describes how one's credences should *change*. Regularity, in contrast, is synchronic. For this reason, regularity should really be compared to a synchronic norm, such as QWAC.

Second, regularity is inconsistent with probabilism when there are uncountably many disjoint events, as for example, when one is uncertain about the value of a physical constant that might be represented by any real number. Because such cases are ubiquitous in science, Bayesian statisticians retain probabilism but often relax regularity to require only that an experimenter's credences are representable by a probability measure that has *full support*, i.e., that it assigns positive probability to all non-empty *open* sets.[22]

In contrast, both DCC and QWAC are consistent with probabilism, no matter the size of the algebra of events. That implies both DCC and QWAC are consistent with assigning probability zero to some hypotheses. But is a scientist really "cooperative" if she assigns probability zero to some live hypotheses? We think so, as long as her degrees of belief have full support on the set of live hypotheses. But we will not defend that claim here.

One might object that we purport to derive a synchronic norm (QWAC) from a diachronic one (DCC), and such an argument could not be valid. Our response has three parts: (i) our argument relies on *two* diachronic norms, (ii) those two norms are "diachronic" in different ways, and (iii) synchronic norms can be derived from an appropriate combination of diachronic ones. In fact, our argumentative strategy is analogous to defenses of regularity (or the weaker full support condition) that are recognized as valid.[23]

For instance, consider the following defense of regularity. Suppose (pace our thesis) a rational agent has probabilistic credences that are updated by conditionalization. Then regularity (or the weaker full support variant) is necessary to guarantee that one's posterior credences concentrate on the true hypothesis.[24] Here, a synchronic norm (regularity) is derived from a norm

[22]Still other philosophers retain regularity by relaxing probabilism and arguing that "open-mindedness" *requires* that one's credences are representable by hyperreal numbers, which violate an Archimedean condition. See [Easwaran, 2014] for a summary and criticism of those arguments.

[23]Of course, the premises of such defenses are controversial.

[24]See [Freedman, 1963], which shows that having a prior with full support is also sufficient if the parameter space is finite dimensional and insufficient in the infinite dimensional case. For an up-to-date bibliography of "Bayesian consistency" results, see [Shalizi, 2020].

requiring asymptotic reliability (concentration of one's posterior) via a norm (conditionalization) that specifies how one's credences change with *each* new piece of evidence. Analogously, we derive a synchronic norm (QWAC) from a norm requiring asymptotic reliability (cooperativeness) via the Assumption of the previous section, which specifies how one's credal ordering changes with *each* new piece of evidence

3 Further Objections and Replies

Norms often conflict, and so norms of individual rationality likely conflict with norms of collective science [Mayo-Wilson et al., 2011]. Thus, one might object that our appeal to both types of norms to justify public probabilism is illegitimate, as contradictory norms could justify any thesis whatsoever.

Such an objection proves too much, as it applies to any normative thesis about a juror's duties, a doctor's obligations, a sporting referee's responsibilities, and so on. For instance, Bayesian decision theory entails one should never turn down free information [Good, 1967], but jurors are required to turn down free evidence if it is inadmissible.[25] However, despite the apparent conflict between norms for jurors and norms of individual rationality, no one thinks that "any thesis whatsoever" about jurors' obligations can be justified by the two types of norms.

Where does the objection go wrong? To answer that question, note there are at least two ways the objection can be disambiguated. First, our critics might claim that a scientist's duties might conflict with her duties as a private citizen, in the same way a criminal lawyer's obligation to defend a client she knows to be guilty might conflict with her duty, qua private citizen, to report known criminal activity. Of course, we agree that private and professional obligations often conflict. But notice the conflicts in the lawyer's private and professional duties do nothing to challenge the thesis that, qua lawyer, she is obliged to defend her client. Similarly, the fact that a scientist has conflicting private and professional duties does little to refute our thesis that, qua scientist, she should have probabilistic credences.

Second, the objector might continue that our defense of public probabilism is importantly disanalogous from the examples involving lawyers and jurors. Lawyers' professional duties qua lawyers *override* their obligations qua private citizens; similar remarks apply to jurors. Conflict is avoided by specifying a priority of duties. In contrast, a defense of public probabilism might, we have claimed, require appeals to both norms of collective science

[25] Thanks to Blinded for review for this example.

and norms of individual rationality; if one of those types of norms outranks the other in cases of conflict, one may be unable to use both types in defense of public probabilism.

We think this objection gains plausibility by trading on ambiguity. It may be plausible that *some* norms of collective science conflict with *some* norms of individual rationality. But thusfar, we have little reason to suspect the conflicting norms are the ones required to justify public probabilism.

4 Conclusions and Future Work

We have argued that, although rational private credences may violate an Archimedean condition, a scientist's public credences may not. But our argument is only a small step towards justifying public probabilism. Future work should investigate to what extent collective norms of science can justify other "axioms" for credence that (a) are necessary for probabilism but (b) may be poorly justified by norms of individual rationality. Similar work should investigate the possibility of collectivist foundations for conditionalization, the central diachronic norm of Bayesianism.

Finally, we think our framework opens up potential avenues for justifying other parts of Bayesian statistical practice that have received less attention from philosophers. As we mentioned in the introduction, Bayesian philosophers have yet to defend the various conventional prior distributions and loss functions that Bayesian statisticians adopt in practice. From the traditional individualistic standpoint, it is potentially irrational to analyze one's data using a prior or loss function that fails to represent one's credences or preferences. Are there, for example, collectivist defenses of conjugate priors? Priors good frequentist properties? And so on. Similarly, can appeal to norms of collective science help justify the computational techniques (e.g., Markov-Chain Monte Carlo methods) that Bayesians use to approximate their posterior distributions when an exact calculation is impossible?

References

Editorial Board. Peer-review policy | Nature Research, October 2020. URL https://www.nature.com/nature-research/editorial-policies/peer-review#the-review-process.

L. Jonathan Cohen. *An essay on acceptance and belief*. Oxford: Clarendon Press, 1992.

Kenny Easwaran. Regularity and hyperreal credences. *Philosophical Review*, 123(1):1–41, 2014.

Anthony William Fairbank Edwards. *Likelihood*. CUP Archive, 1984.

Lina Eriksson and Alan Hájek. What are degrees of belief? *Studia Logica*, 86(2):183–213, 2007. URL http://link.springer.com/article/10.1007/s11225-007-9059-4.

Peter C. Fishburn. The axioms of subjective probability. *Statistical Science*, 1(3):335–358, 1986. URL http://projecteuclid.org/euclid.ss/1177013611.

Will Fleisher. Rational endorsement. *Philosophical Studies*, 175(10):2649–2675, 2018.

David A. Freedman. On the asymptotic behavior of Bayes' estimates in the discrete case. *The Annals of Mathematical Statistics*, pages 1386–1403, 1963. Publisher: JSTOR.

Andrew Gelman and Cosma Rohilla Shalizi. Philosophy and the practice of Bayesian statistics. *British Journal of Mathematical and Statistical Psychology*, 66(1):8–38, 2013.

Irving John Good. On the principle of total evidence. *The British Journal for the Philosophy of Science*, 17(4):319–321, 1967.

Remco Heesen. Why the reward structure of science makes reproducibility problems inevitable. *The journal of philosophy*, 115(12):661–674, 2018.

James M. Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4):575–603, 1998.

Philip Kitcher. The Division of Cognitive Labor. *The Journal of Philosophy*, 87(1):5–22, 1990.

Philip Kitcher. *The Advancement of Science: Science Without Legend, Objectivity Without Illusions*. Oxford University Press, USA, 1995.

David H. Krantz, R. Duncan Luce, Patrick Suppes, and Amos Tversky. *Foundations of Measurement Volume I: Additive and Polynomial Representations*. Dover Publications, Mineola, NY, December 2006.

Isaac Levi. *Gambling with Truth*. MIT Press, Cambridge, 1967.

Isaac Levi. *The enterprise of knowledge: An essay on knowledge, credal probability, and chance*. MIT press, 1983.

Patrick Maher. *Betting on theories*. Cambridge University Press, 1993.

Deborah G. Mayo. *Statistical inference as severe testing*. Cambridge: Cambridge University Press, 2018.

Conor Mayo-Wilson, Kevin J. Zollman, and David Danks. The Independence Thesis: When Individual and Social Epistemology Diverge. *Philosophy of Science*, 78(4):653–677, 2011.

Robert K. Merton. *The sociology of science: Theoretical and empirical investigations*. University of Chicago press, 1973.

Brian A. Nosek, Charles R. Ebersole, Alexander C. DeHaven, and David T. Mellor. The preregistration revolution. *Proceedings of the National Academy of Sciences*, 115(11):2600–2606, 2018.

Cailin O'Connor and Justin Bruner. Dynamics and diversity in epistemic communities. *Erkenntnis*, 84(1):101–119, 2019. Publisher: Springer.

Samir Okasha. On the interpretation of decision theory. *Economics & Philosophy*, 32(3):409–433, 2016.

Richard Pettigrew. *Accuracy and the Laws of Credence*. Oxford University Press, 2016. ISBN 978-0-19-873271-6.

Hanna Rosin. Rock of Ages, Ages of Rock (Published 2007). *The New York Times*, November 2007. ISSN 0362-4331. URL https://www.nytimes.com/2007/11/25/magazine/25wwln-geologists-t.html.

Richard Royall. *Statistical evidence: a likelihood paradigm*, volume 71. Chapman & Hall/CRC, 1997.

Leonard J. Savage. *The foundation of statistics*. Dover publications, 1972.

Teddy Seidenfeld, Mark J. Schervish, and Joseph B. Kadane. Coherent choice functions under uncertainty. *Synthese*, 172(1):157, March 2009.

Amartya K. Sen. Choice functions and revealed preference. *The Review of Economic Studies*, 38(3):307–317, 1971.

Cosma Rohilla Shalizi. Frequentist Consistency of Bayesian Procedures, 2020. URL http://bactra.org/notebooks/bayesian-consistency.html.

Michael Strevens. The Role of the Priority Rule in Science. *The Journal of philosophy*, 100(2):55–79, 2003.

Bas C. Van Fraassen. *The scientific image*. Oxford University Press, USA, 1980.

Susan Vineberg. Dutch Book Arguments. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2016 edition, 2016. URL https://plato.stanford.edu/archives/spr2016/entries/dutch-book/.

Peter Walley. *Statistical reasoning with imprecise probabilities*. Number 42 in Monographs on Statistics and Applied Probability. Chapman & Hall, New York, 1991.

Kevin JS Zollman. The credit economy and the economic rationality of science. *The Journal of Philosophy*, 115(1):5–33, 2018.