

THE REVIEW OF SYMBOLIC LOGIC
Volume 0, Number 0, Month 2009

MERGING OF OPINIONS AND PROBABILITY KINEMATICS

SIMON M. HUTTEGGER

University of California at Irvine

Abstract. We explore the question of whether sustained rational disagreement is possible from a broadly Bayesian perspective. The setting is one where agents update on the same information, with special attention being given to the case of uncertain information. The classical merging of opinions theorem of Blackwell and Dubins shows when updated beliefs come and stay closer for Bayesian conditioning. We extend this result to a type of Jeffrey conditioning where agents update on evidence that is uncertain but solid (hard Jeffrey shifts). However, merging of beliefs does not generally hold for Jeffrey conditioning on evidence that is fluid (soft Jeffrey shifts, Field shifts). Several theorems on the asymptotic behavior of subjective probabilities are proven. Taken together they show that while a consensus nearly always emerges in important special cases, sustained rational disagreement can be expected in many other situations.

The criteria incorporated in the personalistic view do not guarantee agreement on all questions among all honest and freely communicating people, even in principle.

L. J. Savage, *Foundations of Statistics*

§1. Introduction Agreement appears to be part of what it is to respond rationally to new information. If you and I change our beliefs rationally when confronted with the same evidence, shouldn't we reach a consensus? This idea belongs to what Richard Jeffrey calls 'dogmatic epistemology'—it is an expectation "rooted in cases where simple acceptance and rejection are appropriate attitudes toward hypotheses" (Jeffrey, 1987, 391). Suppose we both use first-order logic and have the same evidence in the sense that we agree that the premises of a certain argument are true. If you deductively infer a conclusion from these premises, then I cannot rationally disagree with you about the truth value of the conclusion.

Yet there is disagreement on all kinds of subject matters—scientific, religious, ethical, aesthetic, to name just a few. Does disagreement imply that not everyone's beliefs are justified, that is, that someone must be irrational?

This issue is related to the current debate on peer disagreement in epistemology (e.g. Gutting, 1982; Kelly, 2005; Christensen, 2007, 2009; Elga, 2007). The central question of peer disagreement is how, or whether, you should change your beliefs after learning that they don't agree with the beliefs of someone you consider to be a peer in all relevant respects. Intuitions pull in two directions. On the one side, since the other one is a peer, you might want to revise your beliefs so that they are closer to her's. This is the conciliatory view. On the other side, how could a peer possibly disagree with you in the first place, provided that you both have the same information? This leads to steadfast views, which tell you to stick to your beliefs. The main question here is how one's judgements of others should affect one's own degrees of belief.

The social aspect of epistemic judgements also plays an important role in an older philosophical tradition (see, e.g., Dewey, 1929; Levi, 1980). Of particular importance to the point of view developed in the present essay is C. S. Peirce's *The Fixation of Belief* (Peirce, 1997). Pierce discusses at length the fixation of opinions in a community, emphasizing the advantages of the scientific method which resolves interpersonal disagreements by reference to shared experimental evidence. In situations where it is possible to obtain additional experimental evidence, this provides an alternative to the setup discussed in the peer disagreement literature.

The purpose of this essay is to develop the idea of reaching a consensus by obtaining more evidence with the help of Bayesian epistemology. Bayesianism provides a coherent and precise conceptual framework for studying questions in epistemology, as long as one is prepared to think that subjective degrees of belief are probabilities. I will assume this throughout what follows; but let me add that even if one has doubts concerning the basic principles of Bayesianism, it is worthwhile to apply the apparatus of probability theory as an approximate model for subjective beliefs, for it allows one to get fairly clear answers to general philosophical questions. It is a reasonable hope that answers thus obtained will be qualitatively similar for other theories of belief and belief change.

Bayesians usually conceive of belief revision in terms of conditionalization: new information is incorporated into one's prior degrees of belief by conditioning on it. In this learning situation problems of disagreement do not arise for certain varieties of objective Bayesianism.¹ Suppose, for instance, that there is a uniquely rational prior probability, which can be unearthed by applying methods like the principle of indifference. By starting with the same prior, two conditionalizers will always agree as long as they get the same information.²

But objective Bayesianism runs into a host of objections,³ which can be avoided by adopting a subjective interpretation of Bayesian principles. Subjective Bayesians deny the existence of a uniquely rational prior. Agents are allowed to have (perhaps

¹ Jaynes (2003) is a well known example. Carnap's early work on inductive logic also belongs in this camp (Carnap, 1950). Laplace and plausibly Bayes himself can also be counted as objective Bayesians.

² An alternative to conditioning on the same information is developed by Aumann (1976), who demonstrated that two agents with the same prior will have the same posterior if their posteriors are common knowledge. The agents agree regardless of whether they share their evidence or their posteriors. But the posterior they come to agree depends on what they are sharing.

³ See Howson & Urbach (1993) for an overview and references to the relevant literature.

wildly) different initial degrees of belief. In general, no set of initial beliefs is more, or less, justified than another. This implies that conditioning on the same information can lead two agents to have different posterior beliefs. At this level of analysis it is therefore clear that subjective Bayesianism allows for rational disagreement. This is not too surprising, since subjective Bayesianism is a theory of consistency: it does not tell you what to believe, it only tells you when your beliefs and your updated beliefs are internally consistent. So, while belief formation is not just as you like it, a lot is up to you.

Is this all Bayesians have to say about disagreement? Suppose it were. Then subjective Bayesianism faces a crucial problem, for it seems that subjective Bayesians have no account at all for the connection between rational learning and reaching a consensus. One could conclude from this that subjective Bayesians cannot give good reasons for the rationality of science, as e.g. Chalmers (1999) seems to suggest. One could even try to argue that this provides a good reason to return to objective Bayesianism, despite its shortcomings.

There is, however, a Peircean response to this charge that has been noted by many authors and is explored more fully in this essay. The basic thought is that disagreement often turns out to be transient and disappears as one gets more information. In other words, as more evidence becomes available a consensus may emerge. The underlying empiricist credo is that experience trumps any initial belief state; diverging opinions are just a sign that not enough evidence has accumulated yet.

There is a mathematically clearcut statement of this idea. It is known as ‘merging of opinions’. In a seminal paper, Blackwell & Dubins (1962) showed that merging of opinions holds for Bayesian conditioning.⁴ The Blackwell-Dubins theorem specifies under what conditions two conditional probability measures will become and remain close. If the two probability measures are viewed as prior degrees of belief, this means that agents will expect their conditional beliefs to be approximately the same in the long run. The result holds whenever a particular requirement on prior probabilities, called absolute continuity, is met. If this condition is not met, then the conditional beliefs of two agents may not merge. I explain this idea and discuss its philosophical significance in §2. and §3.

The main difference between this approach to consensus formation and the solutions offered in the peer disagreement literature is that the former operates solely based on the basic principle of Bayesianism—coherence—while the latter adds principles over and above coherence. This is not to say that these additional principles become obsolete in the face of results like the Blackwell-Dubins theorem; rather, such results show under what conditions consensus may be reached in the absence of additional principles.

The Blackwell-Dubins theorem is a striking and deep result which solves several questions concerning long run consensus for Bayesian conditioning. There remain open questions, however. One is the following: are there any other factors (besides the ones required in the Blackwell-Dubins theorem) that influence whether opinions merge? An alternative factor could be the type of evidence one gets. Bayesian

⁴ Earlier, Savage (1954) developed a merging of opinions result for the special case of i.i.d. sequences of random variables with a finite parameter space and non-extreme priors over that space.

conditioning assumes that beliefs change by learning a proposition for certain. But learning experiences need not be like that at all. Evidential inputs may be, and very often are, uncertain. A well-known Bayesian model of belief change for uncertain evidence—and the one considered in the present essay—is Richard Jeffrey’s *probability kinematics* (Jeffrey, 1957, 1965, 1968), which I shall review in §4. While probability kinematics does allow for evidence to be uncertain, it is still close enough to Bayesian conditioning that it might be related to the Blackwell-Dubins theorem.

It seems plausible that uncertain evidence can affect the possibility of reaching a consensus. For the case of probability kinematics, we shall see that this depends on how one understands uncertain evidence. I consider two possibilities: hard Jeffrey shifts, where prior information is ignored; and soft Jeffrey shifts, which take into account an agent’s prior. In a sense, hard Jeffrey shifts, while being uncertain, represent a more solid kind of evidence than soft Jeffrey shifts. They allow to say more explicitly what the uncertain evidence is, while soft Jeffrey shifts can be thought of as interpretations of evidential inputs in the light of one’s prior probabilities.

I show in §5. that probability kinematics with hard Jeffrey shifts leads to merging of opinions under quite natural conditions. This bears on the reasonableness of subjective Bayesianism because hard Jeffrey shifts arguably belong to the kind of evidence, together with propositions, that one might often find in scientific investigations. On the other hand, in §6. I show that soft Jeffrey shifts allow beliefs not to merge even if we impose all the conditions that otherwise entail merging. Thus, uncertain evidence can have a strong impact on whether a consensus is reached. Even if circumstances are very favorable—in particular, if every agent is individually rational—soft kinds of evidence are compatible with persistent disagreement.⁵

Although the question of agreement is not independent of the kind of evidence agents get, other consequences of rational learning from experience are. One example, which I discuss in §7., is the convergence of degrees of belief. I will prove that probability kinematics leads to stable degrees of belief in the long run, provided that learning is rational. This is of course not the same as reaching a long run consensus. But it at least shows that rational learning leads to maximally informed opinions regardless of how uncertain evidential inputs are.

The philosophical significance of these results crucially depends on the Bayesian concept of rational learning. An appropriate framework will be reviewed in §8., where it is shown that a key assumption follows from dynamic coherence. In fact, considerations of dynamic coherence lead to generalizations of several of our results (see §9.).

The big picture that emerges from these considerations is this. From a Bayesian perspective, disagreement is in general consistent with the agents’ individual rationality, modulo assumptions on priors and the structure of evidence; and it is so at each step of a process of observation: from the level of initial probabilities, to the level of finitely many evidential inputs, to the long run. Reaching a consensus is by no means something that is guaranteed. However, there are special circumstances

⁵ This should not be thought of as supporting the steadfast view in the literature on peer disagreement mentioned earlier. Within the framework of soft Jeffrey shifts agents are not required to stick to their guns; rather there are different admissible ways to respond to the same information.

where it virtually is, and these circumstances turn out to be especially salient for fields of inquiry where people often do seem to come to an agreement. The Bayesian approach advocated in this essay explains in addition how certain factors, such as initial beliefs or weak and uncertain kinds of evidence, can lead to sustained but individually rational disagreement. In §10. I try to put some context to these broader issues.

§2. Convergence to the Truth Let's start with a simple example. Adam and Eve observe a sequence of coin flips. Adam believes that the coin is fair. Eve thinks it is biased. After they have observed the *same* finite sequence of coin flips, they find themselves disagreeing about the probability of getting heads on the next trial. A subjective Bayesian would trace their disagreement back to their priors, maintaining that there is no general way to say that one prior is better than another.

But what if Adam and Eve observe more coin flips? It seems quite possible that their degrees of belief will get closer. This is indeed sometimes the case, and under certain regularity conditions both Adam and Eve expect this to happen. I present a precise statement of this result in the next section. In this section I'd like to discuss a first step towards it, which is known as 'convergence to certainty'. Convergence to certainty asserts that Eve and Adam expect to learn the truth about any coin flipping event in the limit.

We adopt the framework of Kalai & Lehrer (1994), which allows us to state our results in a very streamlined way.⁶ Suppose that Ω is a set of atomic events, and \mathfrak{F} is a σ -field of subsets of Ω (that is, \mathfrak{F} is closed under countable unions and complementation). The elements of Ω can be thought of as possible worlds and the members of \mathfrak{F} as propositions. In the foregoing example, Ω is the set of all infinite sequences of coin tosses, and \mathfrak{F} contains all propositions about coin tossing events of interest. Members of \mathfrak{F} describe finite sequences of coin flips, but also limiting properties of sequences of coin flips, such as $\lim_{n \rightarrow \infty} S_n/n = 1/2$ or $\limsup_{n \rightarrow \infty} S_n \geq 0$, where S_n is the total number of heads in the first n flips of the coin.

Let $\mathfrak{E}_1, \mathfrak{E}_2, \dots$ be an infinite sequence of partitions of Ω , and suppose that \mathfrak{E}_{n+1} is a refinement of \mathfrak{E}_n for any n (every element of \mathfrak{E}_{n+1} is a subset of an element of \mathfrak{E}_n). \mathfrak{E}_n represents the information an agent may get about the actual world $\omega \in \Omega$ at time n . For instance, if a coin is flipped infinitely often, the two propositions

⁶ Merging of opinions was also investigated with somewhat different frameworks. Blackwell & Dubins (1962) formulate their result in terms of product probability spaces. In philosophy, the merging of opinion result of Gaifman & Snir (1982) is usually better known than Blackwell-Dubins. Gaifman & Snir work within the framework of a rich logical language.

It is worth saying a bit more about how propositions can be understood here. A standard first order logic—such as the one used by Gaifman & Snir—only admits finite conjunctions and disjunctions. Gaifman & Snir follow Carnap's (1971; 1980) and associate probabilities that are defined over a first order language with probabilities over the set of all models of that language, a sentence being identified with the set of models where it is true. Carnap stipulates that the resulting propositions should be closed under countable operations. This results in propositions that are not expressible in the first order language but can be approximated by its sentences. The resulting framework is essentially the same as the one used in measure theoretic probability.

that (i) the coin comes up heads on the first toss and (ii) the coin comes up tails on the first toss is a partition of the set of all infinite sequences of coin tosses. After observing the first coin toss Adam and Eve learn which proposition contains the actual infinite sequence of coin tosses. The refinement condition says that an agent learns more about the true state of the world as n progresses.

For simplicity, we will assume throughout the paper that each partition \mathfrak{E}_n is finite. If ω is the true state, the agent learns $E_n(\omega) \in \mathfrak{E}_n$ at stage n , where $E_n(\omega)$ denotes the element $E_n \in \mathfrak{E}_n$ with $\omega \in E_n$. Let \mathfrak{F}_n be the field generated by \mathfrak{E}_n (i.e., besides \emptyset and Ω , \mathfrak{F}_n contains all unions of members of \mathfrak{E}_n). In the coin tossing example, \mathfrak{F}_n describes all coin tossing events up to time n . The refinement condition implies that $\mathfrak{F}_n \subset \mathfrak{F}_{n+1}$ for $n = 1, 2, \dots$. In order to guarantee that the sequence of partitions eventually captures every proposition of interest to the agent, we require that the σ -field generated by all the \mathfrak{F}_n coincides with \mathfrak{F} .⁷

For this and the next section we assume that the agent is updating by Bayesian conditioning. That is to say, if ω is the true state of the world, then at each time n the agent's new probability for $A \in \mathfrak{F}$ is given by the conditional probability

$$\mathbb{P}[A|E_n(\omega)] = \frac{\mathbb{P}[A \cap E_n(\omega)]}{\mathbb{P}[E_n(\omega)]}. \quad (1)$$

Here we assume that the probability of each member of any partition has positive prior probability. In what follows we shall use this simplifying assumption, although nothing of substance depends on this condition, and it is dropped in the appendix.⁸ The appendix also explains in more detail that *convergence to certainty* is an immediate consequence of the *martingale convergence theorem* (see Gaifman & Snir (1982) and Schervish & Seidenfeld (1990)). Informally, one can think of a martingale as an infinite sequence of fair gambles (a sequence of gambles for which there is no gambling system you could use to your own advantage). This is expressed by saying that your expected total fortune after the next trial is equal to your present total fortune. On average you neither lose nor win. Martingales are important because they lead to very general laws of large numbers (martingale convergence theorems) that do not depend on the quite stringent conditions required for the standard strong law of large numbers.

In our particular context, the martingale convergence theorem implies that

$$\mathbb{P}[A|E_n(\omega)] \rightarrow \begin{cases} 1 & \text{if } \omega \in A \\ 0 & \text{if } \omega \notin A \end{cases} \quad (2)$$

for almost every ω with regard to the prior probability \mathbb{P} . This means that while there might exist a set of atomic events ω for which the statement of convergence

⁷ Thus we implicitly assume that \mathfrak{F} is generated by a countable sequence of subsets of Ω . In general, the σ -field generated by an arbitrary family of subsets of Ω is the smallest σ -field having all elements of that family as members.

⁸ In our case $\mathbb{P}[A|E_n(\omega)]$ is a version of the conditional probability of A given the σ -field \mathfrak{F}_n . Importantly, the conditional probability of an event given a σ -field is not already a conditional probability measure on all of \mathfrak{F} . If it exists, such a conditional probability measure is a 'regular conditional distribution'. See Schervish & Seidenfeld (1990) for a convergence to certainty result when \mathbb{P} admits a regular conditional distribution. We discuss regular conditional distributions in a bit more detail in the context of the Blackwell-Dubins theorem and in 12.1.

in (2) does not hold, this set has probability zero according to \mathbb{P} . (We will abbreviate this with a.e. (almost everywhere) or a.e. (\mathbb{P}) if we want to emphasize the probability measure with respect to which the almost everywhere statement holds.) Convergence to certainty happens with respect to one's current beliefs. If \mathbb{P} is your prior probability, then consistency requires you to believe with probability one that the occurrence of any event will be revealed in the limit whenever your information increases in the way described above. Because of this result, both Eve and Adam can be sure that their conditional probabilities for any event A will converge either to zero or to one. They are certain that in the long run they will know whether A obtains.

Some thought shows that this is not surprising for events A that are in \mathfrak{F}_n for some n . In this case you will know whether A occurs in the long run. But the convergence of certainty theorem asserts more than just that. Suppose that A is an event in the tail- σ -field, which in our case is given by the intersection of all \mathfrak{G}_n where \mathfrak{G}_n is the σ -field generated by $\mathfrak{E}_{n+1}, \mathfrak{E}_{n+1}, \dots$; for example, A might be the event that $\lim_{n \rightarrow \infty} S_n/n = 1/2$ for infinite sequences of coin flips. In such cases A is an infinitary event whose truth value will not be decided after finitely many observations. However, A is still subject to almost sure asymptotic certainty.

A somewhat different convergence result, which will be important in §7., holds if \mathfrak{F} is not generated by the σ -fields \mathfrak{F}_n . Instead, the σ -field generated by the \mathfrak{F}_n may be a proper subset of \mathfrak{F} . Then there are events in \mathfrak{F} that are not elements of the σ -field generated by the \mathfrak{F}_n . If A is such an event, then convergence to certainty does not hold.⁹ However, the martingale convergence theorem still applies, and it can be used to show that the conditional probabilities of A converge almost surely.

Convergence to certainty can be viewed as a consequence of *dynamic coherence* in the following sense. At each trial n the agent updates her probabilities on an element of the partition \mathfrak{E}_n . This is plain vanilla Bayesian conditioning, which can be justified by a dynamic Dutch book argument in situations where conditioning is the appropriate learning rule.¹⁰ Updating by Bayesian conditioning embeds the sequence of conditional probabilities in the convergence to certainty theorem. Thus, a dynamically coherent agent expects her future degrees of belief to converge to certainty under the appropriate conditions.¹¹

§3. Merging of Opinions Convergence to certainty yields a first pass on merging of opinions.¹² Suppose that Eve's degrees of beliefs are represented by \mathbb{P} and Adam's by \mathbb{Q} , and let $\mathbb{P}_n[A](\omega) = \mathbb{P}[A|E_n(\omega)]$ and $\mathbb{Q}_n[A](\omega) = \mathbb{Q}[A|E_n(\omega)]$. Then $\mathbb{P}_n[A]$ and $\mathbb{Q}_n[A]$ both converge to zero or to one almost surely with respect to the priors \mathbb{P} and \mathbb{Q} , respectively. Now, Eve believes with certainty that $\mathbb{P}_n[A]$

⁹ In this case the indicator of A is not an \mathfrak{F} -measurable random variable. See 12.1. for details.

¹⁰ See Teller (1973) and Skyrms (1987a,b, 1990). Bayesian conditioning may also be justified other arguments; see, e.g., Greaves & Wallace (2006), Leitgeb & Pettigrew (2010), Easwaran (2013) and Huttegger (2013).

¹¹ Schervish & Seidenfeld (1990) present an alternative view on convergence to certainty where the agent need not be dynamically coherent but chooses from a set of posterior probabilities. If this set has some specific properties, convergence to certainty holds.

¹² As noted by, e.g., Earman (1992).

and $\mathbb{Q}_n[A]$ will agree in the limit whenever she assigns probability one to any set to which Adam assigns probability one. For then, since $\mathbb{Q}_n[A]$ goes to certainty a.e. (\mathbb{Q}) , Eve also believes with probability one that her's and Adam's conditional probabilities for any event A are the same in the limit.

This result applies to particular events A . But it does not say anything about Eve's and Adam's overall conditional probabilities. The Blackwell-Dubins theorem fills this gap. It asserts that, as Adam and Eve observe more coin tosses and update by Bayesian conditioning, their degrees of belief will become close *uniformly* in all events. Moreover, the Blackwell-Dubins theorem does not require that conditional probabilities converge (as in convergence to certainty). Eve's and Adam's conditional probabilities may get closer even if they don't converge.

A precise formulation of this result requires us, in the first place, to say more about conditional probabilities. In §2. conditional probabilities were taken with respect to one event A . It is well known that this need not give rise to the existence of a countably additive conditional probability measure for each ω . If such a measure exists, then the prior admits a *regular conditional distribution* on \mathfrak{F} (for a definition, see 12.1.). In our case, a regular conditional distribution always exists since the partition \mathfrak{E}_n is assumed to be finite.¹³ Whenever $E \in \mathfrak{E}_n$ has positive probability, the regular conditional distribution $\mathbb{P}[\cdot|\mathfrak{F}_n](\omega)$ is equal to the conditional probability $\mathbb{P}[\cdot|E]$ if $\omega \in E$. Hence, up to sets of measure zero, we can take conditional probabilities given events as our main locus of interest. However, convergence to certainty for regular conditional distributions as well as merging of opinions hold more generally; in particular, they might hold for continuous random variables. This is further discussed (together with some important qualifications) in 12.1.

We also need to clarify what it means for two probability measures to be close. We follow Blackwell & Dubins (1962) in using the *variational distance*, which takes the least upper bound of the absolute differences of two measures μ and ν over all events in \mathfrak{F} as their distance:¹⁴

$$d(\mu, \nu) = \sup_{A \in \mathfrak{F}} |\mu[A] - \nu[A]|.$$

If μ and ν are ε -close according to d (i.e., $d(\mu, \nu) < \varepsilon$), it follows that $\mu[A]$ and $\nu[A]$ are ε -close for all A . Since we wish to see when two probability measures become close overall, the variational distance seems to be a natural choice of metric.

Merging is now defined as follows: \mathbb{P} is said to *merge to* \mathbb{Q} if for \mathbb{Q} almost every ω

$$d(\mathbb{P}_n(\omega), \mathbb{Q}_n(\omega)) \rightarrow 0$$

as $n \rightarrow \infty$, where $\mathbb{P}_n(\omega) = \mathbb{P}_n[\cdot|E(\omega)]$, $\mathbb{Q}_n(\omega) = \mathbb{Q}_n[\cdot|E(\omega)]$ are the agents' conditional probability measures. Hence, a.e. (\mathbb{Q}) , given any $\varepsilon > 0$ there is an n_0 such that

$$|\mathbb{P}[A|E_n(\omega)] - \mathbb{Q}[A|E_n(\omega)]| < \varepsilon$$

¹³ For each $E \in \mathfrak{E}_n$ and each $A \in \mathfrak{F}$, if $\mathbb{P}[E] > 0$ let $\mathbb{P}[A|\mathfrak{F}_n](\omega) = \mathbb{P}[A|E]$ for $\omega \in E$, and if $\mathbb{P}[E] = 0$ let $\mathbb{P}[A|\mathfrak{F}_n](\omega) = \mathbb{P}[A]$ for $\omega \in E$; see Billingsley (2008, p. 438).

¹⁴ The variational distance induces the so-called *strong topology* on a space of probability measures. Another topology often used for spaces of probability measures is the *weak topology* whose notion of a limit is based on weak convergence. Two probability measures that are close in the strong topology are also close in the weak topology.

for all $n > n_0$ and for all A in \mathfrak{F} . The number n_0 may depend on ε and on ω but not on A .¹⁵ If \mathbb{P} merges to \mathbb{Q} , then Adam believes with probability one that Eve's and his conditional degrees of belief for all propositions A will get arbitrarily close.

Kalai & Lehrer (1994) use an alternative notion of closeness between probability measures. The topology it yields is not the same as the topology of the variational norm. But the two notions of distance are asymptotically equivalent (Kalai & Lehrer, 1994). As a consequence, Kalai and Lehrer's asymptotic results also hold in our framework.

The second element that we need in order to state the Blackwell-Dubins theorem is absolute continuity. \mathbb{Q} is *absolutely continuous relative to* \mathbb{P} (in symbols $\mathbb{Q} \ll \mathbb{P}$) if for all $A \in \mathfrak{F}$

$$\mathbb{Q}(A) > 0 \implies \mathbb{P}(A) > 0.$$

In other words, \mathbb{Q} agrees with \mathbb{P} on the events that are considered virtually impossible (have zero probability) by \mathbb{P} .

Blackwell & Dubins (1962) prove that *if $\mathbb{Q} \ll \mathbb{P}$, then \mathbb{P} merges to \mathbb{Q}* . This says that absolute continuity all but guarantees that conditional probabilities will get and stay closer as more information becomes available: an agent having \mathbb{Q} as personal probability believes with probability one that \mathbb{P} will merge to \mathbb{Q} whenever $\mathbb{Q} \ll \mathbb{P}$.¹⁶

Merging from the point of view of both prior probability measures is an immediate corollary. It requires that \mathbb{P} and \mathbb{Q} be mutually absolutely continuous: for all $A \in \mathfrak{F}$

$$\mathbb{P}(A) = 0 \iff \mathbb{Q}(A) = 0.$$

In this case, \mathbb{P} and \mathbb{Q} will merge a.e. with respect to both measures \mathbb{P} and \mathbb{Q} . In terms of degrees of belief, both agents believe with probability one in merging of opinions.

Merging of opinions has something to offer for both objective and subjective Bayesians (Diaconis & Freedman, 1986). The interpretation we have adopted above is subjective and takes merging to be a limiting result from the point of view of the subjective probability measure \mathbb{Q} . But suppose that \mathbb{Q} is the true probability, or chance, generating a random process. Then merging of opinions means that the chance of an agent's conditional probability measure converging uniformly to the true conditional probabilities is one. We have merging of opinions in chance. Absolute continuity here means that the agent will never be surprised by learning something she thought to be impossible.

What happens if absolute continuity fails? Then degrees of belief need not merge. Consider the following example where Adam's prior \mathbb{P} is not absolutely continuous relative to Eve's prior \mathbb{Q} . Let \mathbb{P} and \mathbb{Q} be two exchangeable probability measures on denumerable sequences of coin flips. Exchangeability means that the probability of any finite sequence of coin flips does not depend on the order in which heads and tails appear. By de Finetti's theorem, \mathbb{P} and \mathbb{Q} can be represented as mixtures of independent binomial trials.¹⁷ That is, there is a prior over real numbers $\theta, 0 \leq \theta \leq 1$, such that trials are i.i.d. given θ . Let \mathbb{P} be the probability that puts prior one on

¹⁵ In special cases of merging n_0 is uniform in ω , for instance in the example of tossing a parametric i.i.d. coin.

¹⁶ The proof for our setup is given in Kalai & Lehrer (1994).

¹⁷ See de Finetti (1937).

$\theta = \frac{1}{2}$ —Eve believes that the coin is fair. Let \mathbb{Q} be the probability that puts prior ε on $\theta = \frac{1}{2}$ and prior $1 - \varepsilon$ on the uniform distribution on the interval $[0, 1]$ —Adam believes with probability ε that the coin is fair and with probability $1 - \varepsilon$ that all biases of the coin are equally probable.

Now $\mathbb{P}[A] > 0$ clearly implies $\mathbb{Q}[A] > 0$, and thus $\mathbb{P} \ll \mathbb{Q}$. But the converse is not true. Just observe that according to \mathbb{P} all infinite sequences with a relative frequency of heads other than $1/2$ have probability zero, but according to \mathbb{Q} they have probability $1 - \varepsilon$. If C is the set of all those sequences, it follows that $\mathbb{P}[C|E_n(\omega)]$ is zero for any evidence E_n and also converges to zero almost surely. On the other hand $\mathbb{Q}[C] > 0$, and for each ω in C , $\mathbb{Q}[C|E_n(\omega)]$ converges to one. Hence, while the two posteriors merge with \mathbb{P} -probability one, they do not merge from the perspective of \mathbb{Q} .

Certain qualifications regarding merging of opinions should be mentioned. Joyce (2010) correctly points out that the agreement established by merging of opinions presupposes certain kinds of agreement at other levels. For instance, agents must agree on the information they get. Mutual absolute continuity assures that the agents agree about probability zero events. This is a substantive assumption because, in general, there is no probability measure on an uncountably infinite space with respect to which every other measure is absolutely continuous. Thus, guaranteeing the absolute continuity requirement of the Blackwell-Dubins theorem is not just a matter of finding a probability measure that will always work.

As noted above, if Adam's prior is not absolutely continuous with respect to Eve's, then their beliefs may not merge. Some critics have remarked that merging of opinions thus begs the question, in particular because absolute continuity might seem to be an unnecessarily strong or even arbitrary requirement (Earman, 1992). That it is not unnecessarily strong follows from the fact that absolute continuity is not just a sufficient, but also a necessary condition for merging.¹⁸ One would have to adopt a notion of merging different from Blackwell-Dubins-Kalai-Lehrer in order to escape this finding.¹⁹

Accepting that, absolute continuity might still seem arbitrary. Why should Adam's prior beliefs be absolutely continuous relative to Eve's? One might hope to derive the absolute continuity requirement from other conditions that seem less whimsical. This appears to be very difficult, however. For instance, one could try to get merging from requiring both agents to be *open-minded*. To evaluate this idea, let's return to the example where Adam and Eve try to estimate the chance of a coin coming up heads. Eve's probability measure \mathbb{Q} assigns probability one to all infinite sequences with the limiting relative frequency of heads being $\frac{1}{2}$, while Adam's probability \mathbb{P} is a mixture of the coin being fair and every other bias of the coin being equally probable. Since \mathbb{P} fails to be absolutely continuous relative to \mathbb{Q} , posteriors don't merge with \mathbb{P} -probability one. Yet both \mathbb{P} and \mathbb{Q} are open minded in the sense that they assign positive probability to every finite initial sequence of observations.

There seems to be no obvious way to replace the absolute continuity condition with something that is *prima facie* more plausible. Where does this leave the significance of the Blackwell-Dubins theorem? First of all, it's important to

¹⁸ This is demonstrated by Kalai & Lehrer (1994, Theorem 2).

¹⁹ One possible way is to use a weaker notion of convergence of probability measures, such as weak convergence. See Diaconis & Freedman (1986) and D'Aristotile et al. (1988).

note that some of the initial implausibility of the absolute continuity requirement disappears if we read the Blackwell-Dubins theorem a bit differently. Suppose that Eve holds beliefs \mathbb{P} . She's wondering whether she will end up with probabilities similar to her conditional probabilities if she starts with a different prior measure. In this situation, Eve might be reluctant to consider an alternative prior that fundamentally changes her actual prior—that is, changes her beliefs about what she thinks has positive probability and zero probability.

Furthermore, from a subjective Bayesian point of view there is no obvious need to derive absolute continuity from other requirements such as open mindedness. The philosophy of subjective Bayesianism allows agents to assert or deny certain assumptions concerning the basic structure of probability spaces. De Finetti's theorem for exchangeable sequences of events provides a good example. Along Laplacian lines, one could view the assumption of exchangeability as a formalization of ignorance that any rational agent has to adopt. The subjective Bayesianism of de Finetti would, on the other hand, maintain that exchangeability is a basic assumption about the probabilistic structure of the events under consideration, and that it is up to the agent to judge whether or not events are exchangeable.

The absolute continuity condition in the Blackwell-Dubins theorem seems to have a similar status—in the sense that it also is not necessary for the coherence of personal probabilities, but is a valid subjective judgement. Like exchangeability, it may or may not hold in particular situations. It's up to the agent to determine whether it holds. If it does, then consistency requires the agent to be certain that beliefs merge. If absolute continuity fails to hold, then beliefs need not merge.

Merging of opinions thus tells us under what conditions to expect consensus in the long run provided that agents update beliefs by conditionalization. This is important for Bayesian philosophy of science, since, presumably, the evidential inputs for or against scientific hypotheses sometimes are propositions. Even if scientists initially disagree, they may reach a consensus after having observed a sufficient amount of evidence. The Blackwell-Dubins theorem implies nothing concerning the speed of convergence, though. There is no guarantee that a consensus is reached in a reasonable time. If the learning situation is given more structure, it can be shown that merging takes place quite rapidly (Howson & Urbach, 1993).

As noted in the previous section, conditionalization can be viewed as an expression of Bayesian individual rationality in specific learning situations. Long run consensus is a consequence of dynamic coherence *under the special assumption of absolute continuity*.²⁰ In this case, group agreement is a consequence of individual rationality. However, if long run consensus fails we may be able to trace it back, not to the individual irrationality of an agent, but to the fact that the agents' initial beliefs did not observe the absolute continuity requirement. If the absolute continuity requirement holds, another possible explanation is that merging takes place very slowly; beliefs would merge if agents were given more evidence, but this might not always be possible. The long run may be too long.

²⁰ As in the case of convergence to certainty, Schervish & Seidenfeld (1990) present a variation of the Blackwell-Dubins theorem that does not rely on conditionalization and thus need not be seen as a consequence of dynamic coherence. All that is required is that agents choose posterior probabilities from sets of probabilities with special properties.

§4. Probability Kinematics We have isolated one feature—a failure of absolute continuity—that allows disagreement in the long run to be consistent with individual rationality. This raises the question whether there are other feature that would also lead to this conclusion.

One candidate suggests itself from considering an implicit assumption of conditioning: conditioning implies that what is learned has probability one, and thus the evidential input is *certain*. This does not just mean that the evidential input gets probability one, but that its complement is identified with the empty set, which in turn renders contemplation of further conditioning on the complement meaningless.

Is evidence always certain? You don't need to be a hard-boiled skeptic in order to consider *uncertain evidence* to be ubiquitous. Adam may be uncertain whether the coin came up heads or tails. A scientist may not be entirely certain as to the outcome of an experiment because the underlying process is only accessible through noisy signals. Thus, the results of an observation or introspection need not lead to singling out a proposition that fully captures what has been learned. In such cases Bayesian conditioning is not an adequate model of learning.

The reason why we are interested in uncertain evidence is because it is another potential source of permanent disagreement. Uncertain evidence may not be powerful enough to bring our beliefs arbitrarily close. The next two sections explore this issue. What we need first is a model for updating beliefs on uncertain evidence. Richard Jeffrey's *probability kinematics* (also known as *Jeffrey conditioning*) is a one canonical model for a specific class of learning situations (see, in particular, Jeffrey, 1965).²¹ Probability kinematics can be applied if the evidential input directly affects only the degrees of beliefs for members of a partition \mathfrak{E}_n (recall that each \mathfrak{E}_n is assumed to be finite). That is to say, the conditional probabilities given members of \mathfrak{E}_n are left intact as the agent shifts her beliefs from \mathbb{P}_{n-1} to \mathbb{P}_n :

$$\mathbb{P}_n[A|E] = \mathbb{P}_{n-1}[A|E] \quad \text{for all } E \text{ in } \mathfrak{E}_n \text{ and all } A \text{ in } \mathfrak{F}. \quad (3)$$

In this case, \mathbb{P}_n is said to *come from* \mathbb{P}_{n-1} *by probability kinematics on* \mathfrak{E}_n . \mathfrak{E}_n is a *sufficient partition* for the two element family of probability measures $\{\mathbb{P}_n, \mathbb{P}_{n-1}\}$ (Blackwell & Girshick, 1954; Diaconis & Zabell, 1982). This implies the following rule for calculating the new probabilities for an event A in \mathfrak{F} :

$$\mathbb{P}_n[A] = \sum_{E \in \mathfrak{E}_n} \mathbb{P}_{n-1}[A|E] P_n[E].$$

The values $P_n[E]$ give the new probabilities for members of the partition. They determine the new probability $\mathbb{P}_n[A]$ of every event A by weighing the old conditional probability of A given the members of the partition with the new probabilities for these and summing up the values. This formula follows immediately from condition (3) and the theorem of total probability. If E is learned for certain, then $P_n[E] = 1$ and $\mathbb{P}_n[A] = \mathbb{P}_{n-1}[A|E]$. Thus, conditioning is a special case of probability kinematics. But the latter allows for learning from uncertain evidence, which can lead to values $P_n[E]$ that are positive for every E in \mathfrak{E}_n . In such a situation, the evidence may not rule out the occurrence of any member of \mathfrak{E}_n .

²¹ Levi (1980) presents another approach to avoid the problematic incorrigibility aspect of conditioning, which does not refer to uncertain evidence.

Suppose that $\mathfrak{E}_1, \mathfrak{E}_2, \dots$ is a sequence of partitions that meets the conditions stated in §2. It can be shown that if \mathbb{P}_n comes from \mathbb{P}_{n-1} by probability kinematics on \mathfrak{E}_n for every $n = 1, 2, \dots$ (where $\mathbb{P}_0 = \mathbb{P}$), then \mathbb{P}_n comes from \mathbb{P} by probability kinematics on \mathfrak{E}_n .²²

$$\mathbb{P}_n[A] = \sum_{E \in \mathfrak{E}_n} \mathbb{P}[A|E]P_n[E]. \quad (4)$$

Probability kinematics is a natural generalization of conditioning for situations where the evidential input only yields uncertain information about a partition. While being applicable to a broader class of learning situations, it preserves many properties of conditioning (Diaconis & Zabell, 1982; Jeffrey, 1988). For the next sections we only consider the case where an agent assumes to undergo learning experiences that yield a sequence P_1, P_2, \dots of probability measures on $(\Omega, \mathfrak{F}_1), (\Omega, \mathfrak{F}_2), \dots$. More to the point, our agent believes with probability one that she will revise her probabilities by performing probability kinematics with P_1, P_2, \dots . Each probability measure P_n is fully determined by attaching probability values to members of \mathfrak{E}_n . In §8. we are going to relax this assumption and allow P_1, P_2, \dots to be random distributions, meaning that an agent can be genuinely uncertain about uncertain future learning experiences.

Two basic assumptions concerning (4), one minor and one major, will play an important role in the following sections. In order to explain the major assumption, it helps to note a consequence of absolute continuity. Let $P_n \ll \mathbb{P}|_{\mathfrak{F}_n}$ denote that P_n is absolutely continuous relative to \mathbb{P} on (Ω, \mathfrak{F}_n) .

(AC) If $P_n \ll \mathbb{P}|_{\mathfrak{F}_n}$, then for every $\varepsilon > 0$ there is a $\delta_n > 0$ such that

$$\mathbb{P}[B] < \delta_n \implies P_n[B] < \varepsilon$$

for all $B \in \mathfrak{F}_n$.²³

For some of our main theorems we require something stronger, namely that this relation holds uniformly in the posteriors P_1, P_2, \dots . Let us say that the sequence P_1, P_2, \dots of probability measures on $(\Omega, \mathfrak{F}_1), (\Omega, \mathfrak{F}_2), \dots$ is *uniformly absolutely continuous* with respect to the measure \mathbb{P} on \mathfrak{F} if

- (i) $P_n \ll \mathbb{P}|_{\mathfrak{F}_n}$ for each n and
- (ii) for every $\varepsilon > 0$ there is a $\delta > 0$ such that for all n

$$\mathbb{P}[B] < \delta \implies P_n[B] < \varepsilon$$

²² Since each \mathfrak{E}_{n+1} is a refinement of \mathfrak{E}_n , $\mathbb{P}_{n+1}[A|E] = \mathbb{P}[A|E]$ for $n = 0, 1, 2, \dots$ and $E \in \mathfrak{E}_{n+1}$. This follows from the fact that probability kinematics leaves the ratios of subsets of members of the partition invariant: $\mathbb{P}_{n+1}[A \cap E]/\mathbb{P}_{n+1}[B \cap E] = \mathbb{P}_n[A \cap E]/\mathbb{P}_n[B \cap E] = \dots = \mathbb{P}[A \cap E]/\mathbb{P}[B \cap E]$; $A \cap E$ and $B \cap E$ are subsets of a member of a partition since each earlier partition is a coarsening of \mathfrak{E}_{n+1} . Thus, coming from \mathbb{P}_n by probability kinematics on \mathfrak{E}_{n+1} is the same as coming from \mathbb{P} by probability kinematics on \mathfrak{E}_{n+1} .

²³ Here is a quick proof: Suppose that $P_n \ll \mathbb{P}|_{\mathfrak{F}_n}$ but that the conclusion is not true. Then there is an ε_0 and a sequence of events F_1, F_2, \dots in \mathfrak{F}_n with $\mathbb{P}[F_m] < \frac{1}{2^m}$ but $P_n[F_m] \geq \varepsilon_0$ for all m . Now, the event $\limsup_m F_m$ has probability zero, $\mathbb{P}[\limsup_m F_m] = 0$ (this follows from the Borel-Cantelli lemma since $\sum_m \mathbb{P}[F_m] < \infty$). On the other hand, $P_n[\limsup_m F_m] \geq \limsup_m P_n[F_m]$ by Fatou's lemma. Hence $P_n[\limsup_m F_m] \geq \varepsilon_0$, which contradicts absolute continuity of P_n relative to \mathbb{P} .

for all $B \in \mathfrak{F}_n$.²⁴

In the following sections we often assume that P_1, P_2, \dots is uniformly absolutely continuous relative to a prior measure \mathbb{P} . What does it mean? The first condition—that $P_n \ll \mathbb{P}|_{\mathcal{F}_n}$ for all n —guarantees that for each event A in \mathfrak{F} and for all n , $\mathbb{P}[A] = 0$ implies $\mathbb{P}_n[A] = 0$. This assertion clearly holds for conditioning (for regular conditional probabilities, it holds almost surely). Thus, we require that events which are assigned probability zero now are expected to have probability zero in the future. The possible learning experiences of an agent are live possibilities from her current point of view.

Condition (ii) says that this also holds as we go to the limit. Fix some $\varepsilon > 0$. From (AC) we know that for each n there is a $\delta_n > 0$ such that, for all $B \in \mathfrak{F}_n$, if $\mathbb{P}[B] < \delta_n$ then $P_n[B] < \varepsilon$. If the δ_n have to be chosen so that $\inf_n \delta_n = 0$, then $\delta_n \rightarrow 0$ as $n \rightarrow \infty$, and it is possible that there exists a sequence of events F_1, F_2, \dots in $\mathfrak{F}_1, \mathfrak{F}_2, \dots$ such that $\mathbb{P}[F_n] \rightarrow 0$ as $n \rightarrow \infty$ while the sequence $P_1[F_1], P_2[F_2], \dots$ is bounded away from zero. That possibility is ruled out by condition (ii), which requires that δ_n can be chosen uniformly in n . In other words, (ii) guarantees that absolute continuity does not get lost in the limit. This seems to be a plausible assumption, for the same reason as absolute continuity is plausible. If you think that events have a probability bounded away from zero as you go to the limit, then this should already be reflected in your prior. Still, one might view it as too restrictive for some cases. For even if we lose absolute continuity in the limit, we don't lose it at any finite time. It seems to be very difficult to say which point of view is more plausible in general. Instead of parsing things out further, let me just note that in the worst case you should take (ii) as an additional restrictive assumption that excludes certain learning processes. Also, as explained below, (ii) can often be replaced with a less restrictive but also somewhat more complicated condition.

The second assumption that we use in the next sections is really just a minor convention. We stipulate that in (4) $0 \cdot \mathbb{P}[A|E] = 0$ whenever $\mathbb{P}[E] = 0$. In this case, it follows from absolute continuity (i) that $P_n[E] = 0$. Thus $\mathbb{P}[A|E]P_n[E]$ is well defined even if $\mathbb{P}[A|E]$ isn't. (Generally speaking, when working with regular conditional distributions this requirement is redundant since then conditional probabilities are defined even if the conditioning event has probability zero.)

§5. Merging for Probability Kinematics Suppose now that there are two prior probability measures \mathbb{P} and \mathbb{Q} which are updated successively by probability kinematics on $\mathfrak{E}_1, \mathfrak{E}_2, \dots$ using the distributions P_1, P_2, \dots and Q_1, Q_2, \dots , respectively. Using (4) this leads to the new probability measures $\mathbb{P}_n, \mathbb{Q}_n$ on \mathfrak{F} for $n \geq 1$.

In the present context we say that \mathbb{P}_n merges to \mathbb{Q}_n if $d(\mathbb{P}_n, \mathbb{Q}_n) \rightarrow 0$. Hence for each $A \in \mathfrak{F}$ $|\mathbb{P}_n[A] - \mathbb{Q}_n[A]|$ goes to 0 as n tends to ∞ . Recall that for the Blackwell-Dubins notion of merging we had a \mathbb{Q} -a.e. qualification. We only drop this here for ease of exposition. If the P_n and Q_n are thought of as random distributions and

²⁴ If P_n is viewed as a random variable as suggested in the previous paragraph, it is enough to require that these conditions hold with probability one; cf. Theorems 9.5. and 9.6.

not as fixed sequences, then we would have to add the qualification. We return to this issue in §9.

It is quite obvious that arbitrary choices of sequences P_1, P_2, \dots and Q_1, Q_2, \dots need not lead to merging. But this is also true for conditioning. Recall that one requirement of the Blackwell-Dubins theorem is that agents condition on the *same* factual evidence. Thus, the important question is whether beliefs merge for probability kinematics whenever P_n and Q_n represent the same uncertain information. But what does it mean to get the same uncertain evidence? An obvious starting point for explicating uncertain evidence are *hard Jeffrey shifts* (Joyce, 2010). A hard Jeffrey shift sets values for P_n regardless of the prior probability \mathbb{P}_{n-1} , and so may destroy any information about the partition that was encoded in the prior.

As an example, consider a measurement instrument that makes noisy observations of a physical process, such as coin flips. Let's call this setup a 'mechanical observer'.²⁵ The probability space (Ω, \mathfrak{F}) represents the set of states and events of the process. At each stage n , the output of the mechanical observer is a probability distribution over the partition \mathfrak{E}_n . The probabilities for members of the partition are determined by repeated previous observations under symmetric conditions in order to specify measurement error. More generally, a hard Jeffrey shift can be viewed as a noisy signal where the noise has the form of a probability distribution over a partition such that the distribution is known to every observer.

In terms of hard Jeffrey shifts, having the same uncertain evidence at stage n means that $P_n = Q_n$ on \mathfrak{F}_n . Suppose, for example, that Adam and Eve are two scientists observing coin flips with the help of a mechanical measurement instrument. They might not feel comfortable approximating their learning process by conditionalization if their measurements are not precise enough. Instead, they plan to update by the same hard Jeffrey shifts at each stage of their experiment. Can they be certain to have similar beliefs after having taken sufficiently many measurements, just as they would if they were conditionalizers?

Equation (4) suggests that it may not be enough to assume absolute continuity for priors in order to guarantee merging. Even if $P_n = Q_n$ for all n , without further constraints it is not clear what the influence of P_n and Q_n on the belief dynamics will be. The constraint that we are going to focus on for now requires \mathbb{P}_n and \mathbb{Q}_n to be stable over time regarding members of partitions. More to the point, we will assume that

$$P_n[F] = P_{n-1}[F] \text{ for all } F \in \mathfrak{F}_{n-1}. \quad (\text{M})$$

Condition (M) implies that $\mathbb{P}_m[F], F \in \mathfrak{F}_{n-1}$ is constant for all $m \geq n - 1$. This means that the information learned regarding a partition is never contradicted by information concerning a finer partition. Within each member of a partition, probabilities can change in the future, but only without altering the probability of that member. This condition describes special kinds of hard Jeffrey shifts where prior information is not completely obliterated.²⁶

Condition (M) is significant because it can be given a justification in terms of dynamic coherence (for a more detailed explanation see §8.). It should be emphasized that (M) presupposes that updating of probabilities by P_n proceeds along

²⁵ See Skyrms (1985).

²⁶ It is reminiscent of the definition of Jeffrey independence in Diaconis & Zabell (1982).

the elements of increasingly fine partitions. Bayesian conditioning is a special case of such a process—and condition (M) clearly holds for conditioning. Certain kinds of mechanical observers may also serve as examples. In general, such a mechanical observer provides agents with probability distributions P_n over \mathfrak{F}_n . This means that in repeated observations of the same type the evidential input of agents consists in a joint probability distribution at each stage n . For instance, if Adam and Eve observe sequences of coin flips, then at each stage n of the experiment their common input is a joint probability P_n over events involving the first n coin flips (events in \mathfrak{F}_n).

Under certain conditions such evidential inputs will not contradict earlier evidential inputs. Consider, for instance, a mechanical observer that outputs a joint probability distribution after each flip of a coin in the following way: Let p_n^h, p_n^t be the probability for heads and tails, respectively, determined by the mechanical observer at trial n . This does not yet determine a joint distribution P_n over \mathfrak{F}_n , which is what we need in order to apply (4). One reasonable assumption is that measurements are independent in the following sense: $P_n(e_1, \dots, e_n) = p_1^{e_1} \cdots p_n^{e_n}$, where e_i stands for h or t . In this case, condition (M) clearly holds since $P_n(e_1, \dots, e_{n-1}) = P_n(e_1, \dots, e_{n-1}, h) + P_n(e_1, \dots, e_{n-1}, t) = p_1^{e_1} \cdots p_{n-1}^{e_{n-1}} (p_n^h + p_n^t) = P_{n-1}(e_1, \dots, e_{n-1})$.

The setup put forward in this section only holds for agents who agree on the joint probability distributions P_n over \mathfrak{F}_n at each stage n as their new evidential input. Any sequence of probabilities such that earlier outputs are marginals of present outputs satisfies (M); that is, (M) holds if for each $F \in \mathfrak{F}_{n-1}$

$$P_{n-1}[F] = \sum P_n[E] \quad (5)$$

where the sum ranges over all $E \in \mathfrak{E}_n$ such that $E \subset F$. This is the sense in which later probabilities do not contradict earlier probabilities.

We are now in a position to state a first theorem for merging of opinions with probability kinematics. Its proof, like the proofs of all other theorems, can be found in the appendix.

THEOREM 5.1. *Suppose that $Q_n = P_n$, that the sequence $Q_n, n = 1, 2, \dots$ is uniformly absolutely continuous relative to \mathbb{Q} , and that $\mathbb{Q} \ll \mathbb{P}$. If condition (M) holds, then $d(\mathbb{P}_n, \mathbb{Q}_n) \rightarrow 0$ as $n \rightarrow \infty$.*

The theorem implicitly assumes that both agents update by probability kinematics on the sequence of partitions. That being so, uncertain evidence in the form of hard Jeffrey shifts leads agents to reach a consensus in the long run whenever suitable conditions obtain.

Aside from the assumption that updates are relative to shared probabilities $Q_n = P_n$, the requirement (M) is fairly liberal. How restricted the class of admissible Q_n is clearly depends on \mathbb{P} and \mathbb{Q} . If \mathbb{P} assigns probability zero to a large number of measurable sets, \mathbb{Q} and Q_n are required to do the same. On the other hand, if the underlying probability space has an appropriate topological structure, and if both \mathbb{P} and \mathbb{Q} are open minded in the sense mentioned earlier—both assign positive probability to any open set—then the posteriors Q_n are not very restricted.

The assumption of uniform absolute continuity is plausible for at least some learning situations, but is not required for proving Theorem 5.1. In order to explain what role this assumption plays in the proof of the theorem, consider the Radon-

Nikodým derivatives $Y_n = dQ_n/d\mathbb{Q}|_{\mathfrak{F}_n}$, which exist by our assumption of absolute continuity.²⁷ Roughly speaking, Y_n describes how much the posterior \mathbb{Q}_n and the prior \mathbb{Q} differ on the partition \mathfrak{E}_n . Assuming that $Q_n = P_n$ and that the sequence $Q_n, n = 1, 2, \dots$ is uniformly absolutely continuous relative to \mathbb{Q} guarantees that the sequence $Y_n, n = 1, 2, \dots$ is uniformly integrable (see Lemma 12.7.), a result that is needed to prove the convergence asserted in Theorem 5.1. But, as is also shown in the appendix (see in particular Lemma 12.8.), for proving the theorem it is actually sufficient that the sequence $d(\mathbb{P}_n, \mathbb{Q}_n)Y_n, n = 1, 2, \dots$ is uniformly integrable. While this may be less restrictive than the uniform integrability of the Radon-Nikodým derivatives, it does not have as straightforward an interpretation as uniform absolute continuity.

The present setup arguably covers the structure of some scientific investigations in an idealized sense. It deals with situations where experience delivers joint distributions over increasingly fine-grained partitions, and where agents agree on those distributions. This might be because the distributions are outputs of special types of mechanical observers. But it could also constitute a consensus reached by agents as to the probability of experiments when errors and noise are taken into account. In any case, agents are assumed to defer to the joint probabilities (in the sense of epistemic deference as described, e.g., in Joyce, 2007).

What happens if condition (M) fails? I hope to relax the restriction that later evidence does not contradict earlier evidence as expressed by (M') in §9. by a similar one. But merging is also feasible for another class of learning situations. The relevant condition is that the agent's learning experiences—more precisely, the rates at which beliefs concerning members of a partition change—are not too extreme. This can again be captured by the Radon-Nikodým derivatives Y_n . If, for all n , $Y_n \leq g$ for some non-negative integrable function g , then the differences between posteriors and priors are not too pronounced.

As an example, consider a sequence of distributions Q_1, Q_2, \dots with uniformly bounded Radon-Nikodým derivatives. That is, there is a constant $K < \infty$ that does not depend on n or E with $Q_n[E]/\mathbb{Q}[E] \leq K$ for all n and $E \in \mathfrak{E}_n$. This means that $Q_n[E]$ can be at most K times $\mathbb{Q}[E]$. This and the more general assumption might be reasonable in situations where prior to an investigation one is confident that the evidence is not going to change one's prior radically. Unlike updates observing the martingale condition (M), Q_n does not need to be equal to the marginal \mathbb{Q}_{n+1} as in (5).

If we substitute this new condition for condition (M), the conclusion of Theorem 5.1. continues to hold. Notice that we only require absolute continuity of posteriors relative to the prior and not uniform absolute continuity.

²⁷ If $Q_n \ll \mathbb{Q}|_{\mathfrak{F}_n}$, then the Radon-Nikodým derivative Y_n exists and is almost surely unique. In our setting, the Radon-Nikodým derivative $dQ_n/d\mathbb{Q}|_{\mathfrak{F}_n}$ has a particularly simple form: for almost every ω

$$\frac{dQ_n}{d\mathbb{Q}|_{\mathfrak{F}_n}}(\omega) = \sum_{E \in \mathfrak{E}_n} I_E(\omega) \frac{Q_n[E]}{\mathbb{Q}[E]}.$$

That is, if $\mathbb{Q}[E] > 0$ then it is basically equal to the ratio $Q_n[E]/\mathbb{Q}[E]$ on E .

THEOREM 5.2. *Suppose that $Q_n = P_n$ and $Q_n \ll \mathbb{Q}|_{\mathfrak{F}_n}$ for all $n = 1, 2, \dots$, and that $\mathbb{Q} \ll \mathbb{P}$. If there is a nonnegative function g integrable with respect to \mathbb{Q} such that $Y_n \leq g$ a.e. (\mathbb{Q}) for all n , then $d(\mathbb{P}_n, \mathbb{Q}_n) \rightarrow 0$ as $n \rightarrow \infty$.*

We have seen in this section that opinions merge for probability kinematics with hard Jeffrey shifts under a variety of conditions. Absolute continuity still plays a central role. By virtue of Dutch book theorems for probability kinematics (Armendt, 1980; Skyrms, 1987b, 1990), merging can be viewed as a consequence of dynamic coherence. But the force of this conclusion depends on how to justify the other assumptions of the theorems. I postpone a discussion of these issues to after having introduced some further results in the next three sections. In these sections, condition (M) will again play a crucial role, so please bear in mind that we will later show that it can be extended and also derived from first principles.

§6. Divergence for Probability Kinematics According to van Fraassen (1980) experience speaks to the conditionalizer “with the voice of an angel”. In the case of hard Jeffrey shifts, experience still has the voice of an angel, but it is received through a noisy channel. The deliverances of experience are not as lucid as being delivered the truth of a proposition. Notwithstanding the noisy character of hard Jeffrey shifts, we saw that one of the most important results for conditionalization—that opinions merge with increasing information—can be recovered under certain conditions.

However, experience does not always speak with an angel’s voice. People with the same evidential input may be led to different opinions because of interpreting that evidence differently. How they arrive at their opinions is a complicated process, and I don’t attempt to give an account of it here. I’d rather like to introduce one way to model such learning experiences that is particularly congenial to the Bayesian approach. On this account, experience can lead one to shift one’s probability distribution in a way that essentially depends on one’s prior. Consider the following example of such a *soft learning experience*, taken from Joyce (2010). Suppose that the partition \mathfrak{E}_1 consists of four elements, and let the learning experience be given by

$$P_1[E_1] = \frac{1}{5}, P_1[E_2] = \frac{3}{10}, P_1[E_3] = \frac{1}{2}, P_1[E_4] = 0.$$

If the agent updates her prior measure \mathbb{P} by probability kinematics on \mathfrak{E}_1 using these probabilities, the result of her learning experience is a hard Jeffrey shift. Now examine a learning experience that is given by

$$P_1[E_1] = 2 \cdot \mathbb{P}[E_1], P_1[E_2] = \frac{1}{2} \cdot \mathbb{P}[E_2], P_1[E_3] = 5 \cdot \mathbb{P}[E_3], P_1[E_4] = 0 \cdot \mathbb{P}[E_4].$$

If the agent’s prior probabilities are $\mathbb{P}[E_1] = \frac{1}{10}, \mathbb{P}[E_2] = \frac{3}{5}, \mathbb{P}[E_3] = \frac{1}{10}, \mathbb{P}[E_4] = \frac{1}{5}$, then probability kinematics will lead to the same result as in the previous example. But there is a profound difference between the two: in the second example, the agent’s new probabilities for elements of the partition depend on her prior probabilities. The learning experience only specifies the ratio of new probability to old probability $P_1[E]/\mathbb{P}[E]$ for each member E of \mathfrak{E}_1 . This is a Bayesian approach to a situation where an agent’s belief system (as given by her prior) influences how she interprets an evidential input. Changing the priors in the second example leads

to degrees of belief different from the ones in the first example, where the prior has no influence on the result of learning. For probability kinematics, soft learning experiences are, not surprisingly, called ‘soft Jeffrey shifts’ (Joyce, 2010). The main difference between hard Jeffrey shifts and soft Jeffrey shifts is that the former, in general, nullifies all information that is contained in the agents prior, while the latter preserves some of that information.

Do beliefs merge when agents have the same soft uncertain evidence? We are going to see that this need not be the case. If Adam and Eve start with different (but mutually absolutely continuous) prior probabilities for infinite sequences of coin flips, and if both observe principle (M) as well as undergo the same soft Jeffrey shifts, their posterior degrees of beliefs may not get close to each other in the long run. Before we are in a position to formulate this result and discuss its consequences, we need to examine soft Jeffrey shifts in more detail.

Soft Jeffrey shifts play an important role in studying the problem of the non-commutativity of probability kinematics (Diaconis & Zabell, 1982; Wagner, 2002). Commutativity means that an agent ends up with the same posterior probability regardless of the order of successive updates. Bayesian conditioning is commutative, but probability kinematics need not be. This has caused some concerns as to whether probability kinematics is a rational way to update beliefs (Döring, 1999; Lange, 2000; Kelly, 2008). I agree with Wagner (2002) and Joyce (2010) that these concerns are misguided. As Joyce (2010) points out, probability kinematics is non-commutative exactly when it should be; namely, when belief revision destroys information obtained in previous updates. For hard Jeffrey shifts, this idea was developed by Diaconis & Zabell (1982). For soft Jeffrey shifts, Wagner (2002), extending results by Field (1978) and Jeffrey (1988), shows that with an appropriate understanding of what it is to get the same uncertain evidence, probability kinematics is commutative.

Wagner’s argument is based on a numerical measure of what is learned from new evidence. This measure is obtained by factoring out the prior. For conditioning, the suitable measure is given by the likelihood ratio (Good, 1950, 1983). *Bayes factors* are a more general measure. Consider two events E, F in the partition \mathfrak{E}_n . The Bayes factor of E and F is given by

$$\mathfrak{B}_{\mathbb{P}_n, \mathbb{P}_{n-1}}(E, F) = \frac{\mathbb{P}_n[E]}{\mathbb{P}_n[F]} \bigg/ \frac{\mathbb{P}_{n-1}[E]}{\mathbb{P}_{n-1}[F]}.$$

For simplicity I will assume throughout this section that $\mathbb{P}_n[E], \mathbb{P}_{n-1}[E] > 0$ for all $E \in \mathfrak{E}_n$ and all n . Appropriate definitions could be used to deal with other cases.

Being ratios of new-to-old odds, Bayes factors are generalizations of likelihood ratios. They can also be formulated in terms of *relevance quotients*. The relevance quotient for E is

$$\mathfrak{R}_{\mathbb{P}_n, \mathbb{P}_{n-1}}(E) = \frac{\mathbb{P}_n[E]}{\mathbb{P}_{n-1}[E]}.$$

Thus, the Bayes factor of E and F is

$$\mathfrak{B}_{\mathbb{P}_n, \mathbb{P}_{n-1}}(E, F) = \frac{\mathfrak{R}_{\mathbb{P}_n, \mathbb{P}_{n-1}}(E)}{\mathfrak{R}_{\mathbb{P}_n, \mathbb{P}_{n-1}}(F)}.$$

A soft Jeffrey shift specifies for each $E \in \mathfrak{E}_n$ the ratio of $\mathbb{P}_n[E]$ and $\mathbb{P}_{n-1}[E]$ up to multiplication by a positive constant. That is, it sets $\mathfrak{R}_{\mathbb{P}_n, \mathbb{P}_{n-1}}(E) = c_n \beta(E)$.

This means that the new probability for E depends on its old probability: $\mathbb{P}_n[E] = c_n \beta(E) \mathbb{P}_{n-1}[E]$. Information contained in the prior is not lost when moving to the new probability. Hartry Field used soft Jeffrey shifts for his reformulation of probability kinematics (Field, 1978). The numbers $\beta(E)$ specify a *Field shift*. For any Field shift there is a corresponding soft Jeffrey shift with

$$P_n[E] = \frac{\beta(E) \mathbb{P}_{n-1}[E]}{\sum_{E \in \mathfrak{E}_n} \beta(E) \mathbb{P}_{n-1}[E]}.^{28}$$

The Field shift given by $\beta(E)$ renders successive belief revisions commutative. For probability kinematics, this leads to criteria for commutativity in terms of certain Bayes factor identities (Wagner, 2002). Suppose that an agent revises her beliefs by probability kinematics twice. Consider the Bayes factors for the first shift in a given order of revisions. If the order is reversed, these Bayes factors have to be identical to the Bayes factors for the second shift in the reversed sequence. This essentially means that probability kinematics will be commutative if the two learning events are the same up to the agent's probability prior to each learning event.

The relationship between commutativity and Bayes factor identities indicates that identical Bayes factors express, in a sense, the information provided by uncertain evidence that is not already contained in the prior. This arguably supplies us with another notion of getting the same uncertain evidence. Let one agent change her degrees of belief from \mathbb{P}_{n-1} to \mathbb{P}_n and another agent from \mathbb{Q}_{n-1} to \mathbb{Q}_n , where both belief revisions are by probability kinematics on \mathfrak{E}_n . We now require that

$$\mathfrak{B}_{\mathbb{P}_n, \mathbb{P}_{n-1}}(E, F) = \mathfrak{B}_{\mathbb{Q}_n, \mathbb{Q}_{n-1}}(E, F) \quad \text{for all } E, F \in \mathfrak{E}_n \quad (6)$$

whenever the agents get the same evidential input.

As explained earlier in this section, this concept of uncertain evidence is clearly different from hard Jeffrey shifts, since responses to learning experiences may depend on an agent's prior. This has important consequences for merging. The next theorem says that having the same Bayes factors is consistent with divergence of opinions. An agent can think she will have the same uncertain evidence (soft Jeffrey shifts) as another agent without believing in merging of opinions. In order to keep its proof manageable, we will make some simplifying and rather innocuous structural assumptions.

N1. We require the following:

If $F \in \mathfrak{E}_n$ then there are $E_1, E_2 \in \mathfrak{E}_{n+1}$ such that $F = E_1 \cup E_2$

for $n = 0, 1, 2, \dots$.²⁹ This means that for the first learning experience Ω is partitioned into two sets E and \bar{E} . For the next learning experience, each of these sets is again partitioned into two sets; and so on. An example of such a sequence of partitions are the partitions generated by tossing a coin infinitely often.

N2. Next we require that the prior probabilities \mathbb{P} and \mathbb{Q} satisfy

$$0 < \mathbb{P}[E], \mathbb{Q}[E] < 1, \mathbb{P}[E] \neq \mathbb{Q}[E].$$

for $\mathfrak{E}_1 = \{E, \bar{E}\}$. This says that priors disagree about the probability of E .

²⁸ Note that the the constant c_n cancels out for the Jeffrey shift.

²⁹ Thus $E_1, E_2 \neq \emptyset$ and $E_1 \cap E_2 = \emptyset$. Also, \mathfrak{E}_0 is the trivial partition $\{\Omega\}$.

N3. The third assumption extends (N2). By (N1), whenever $F \in \mathfrak{E}_n$ there are $E_1, E_2 \in \mathfrak{E}_{n+1}$ such that $F = E_1 \cup E_2$. We now require that, for all n ,

$$\mathbb{P}[E_1|F] \neq \mathbb{Q}[E_1|F]. \quad (7)$$

This is compatible with conditional probabilities being arbitrarily close.

Furthermore, recall our general assumption in this section that the new probabilities P_n, Q_n for members of the partition \mathfrak{E}_n are all positive. These assumptions can now be used to show, first of all, that there exist sequences of new probabilities that satisfy both (M) and (6). The second part of the next theorem then asserts that this may lead to diverging opinions.

THEOREM 6.3. *Suppose that (N1) holds.*

1. *If (N2) and (N3) hold, then there exist sequences of probability measures P_1, P_2, \dots and Q_1, Q_2, \dots on the sequence $(\Omega, \mathfrak{F}_1), (\Omega, \mathfrak{F}_2), \dots$ satisfying condition (M) and (6).*
2. *If the two sequences P_1, P_2, \dots and Q_1, Q_2, \dots satisfy both (M) and (6), then they can be chosen in such a way that, for some $\varepsilon > 0$, $d(\mathbb{P}_n, \mathbb{Q}_n) \geq \varepsilon$ holds for all $n = 1, 2, \dots$*

Earlier we have identified absolute continuity as one necessary element of merging for conditionalization as well as for probability kinematics with hard Jeffrey shifts. For the latter, we explicitly introduced (M) as a condition governing belief change. Theorem 6.3. says that, under the same conditions, uncertain evidence in the sense of soft Jeffrey shifts need not lead to merging of beliefs (in Theorem 6.3. the measures \mathbb{P} and \mathbb{Q} may be absolutely continuous relative to each other). This allows us to conclude that, besides absolute continuity, *the nature and structure of evidence is another determining factor for merging of opinions.*

The proof of Theorem 6.3. is driven by the fact that Bayes factors do not uniquely specify posterior probabilities. Bayes factors only create overall constraints for posterior distributions. This allows agents to interpret the same evidential input in different ways. An interpretation consists in choosing a posterior that is within the constraints given by the Bayes factor identities (6). Theorem 6.3. then shows that disagreement can persist indefinitely among agents with the same uncertain learning experiences even though each agent is individually rational.

We note that we did not identify the Bayes factors with the uncertain evidence an agent gets. We only require that whenever agents get the same uncertain evidence, they will have the same Bayes factors. Bayes factors only partially explicate the type of uncertain evidence. This is in line with Bradley (2005) who argues that a Bayes factor is not identical to an agent's evidential input; rather, it's the agent's reaction to it. Bayes factors do not have the same status as propositions and cannot in general be used by others to update on the same information.³⁰

The deep philosophical problem underlying this and the previous section is the question of how to capture evidential inputs other than factual propositions. This

³⁰ However, both Jeffrey (1992, p. 7-9) and Bradley (2005) observe that if you trust someone's response to experience without sharing her prior, you could apply her Bayes factors to move to a posterior by probability kinematics, again without asserting that the Bayes factor *is* the evidence.

problem becomes especially pressing in the context of merging, which requires us to say rather precisely what it means for two agents to have the same uncertain learning experiences. Both approaches we’ve been considering—soft and hard Jeffrey shifts—are legitimate because they capture different concepts of uncertain evidence. There might well be other concepts of uncertain evidence that may be worth studying with regard to merging of beliefs.

§7. Maximally Informed Opinions There is a result analogous to convergence to certainty for learning from uncertain evidence. Of course, one cannot expect convergence to *certainty* in this case. This point can be explained with the help of a result by Diaconis & Zabell (1982), which tells us when probability kinematics (or any kind of probabilistic learning, for that matter) is equivalent to Bayesian conditioning in a bigger probability space.³¹ Suppose that an agent shifts from a prior to a posterior probability measure by probability kinematics on some partition. Provided that prior and posterior are related in a certain way, one can add two elements a and b to the original sample space: a indicates that the agent had the experience she had, and b indicates its absence. Conditioning on a in the bigger probability space yields the same result as the original shift in beliefs.

With some effort this construction can be extended to infinite sequences of belief revisions by probability kinematics, so that the sequence is represented in terms of conditionalization in a big probability space with phenomenal experiences like a and b as elements. Applying the martingale convergence theorem in this new probability space shows that the sequence of conditional probabilities is a martingale that converges almost surely. However, the smallest σ -field describing the infinite sequences of phenomenal events does not capture all events of interest. It only describes the agent’s phenomenal experiences but not the “real” events such as coin flips. This implies that it will not be known with certainty in the limit whether non-phenomenal events A —that is, events that are not in the σ -field of all phenomenal experiences—have occurred.

This argument suggests that beliefs based on uncertain evidence may converge—the agent’s degrees of belief might settle down not at certainty but at some in between value. Without further constraints on probability kinematics, this need not happen because probabilities for elements of a partition are not restricted and can fluctuate forever. However, this is impossible if condition (M) holds. The following theorem makes this idea precise:

THEOREM 7.4. *Suppose that the sequence $P_n, n = 1, 2, \dots$ is uniformly absolutely continuous with respect to \mathbb{P} . Let A be an element of \mathfrak{F} . If condition (M) holds, then $\mathbb{P}_n[A]$ converges to a limit as $n \rightarrow \infty$ for every $A \in \mathfrak{F}$.*

Theorem 5.2. states that opinions can merge even if condition (M) fails, so long as the agent’s learning experiences are not too extreme. That learning experiences are not too extreme does not guarantee convergence of degrees of belief, however. It is not difficult to see that, for instance, the uniform boundedness of Radon-Nikodým derivatives for the posteriors relative to the prior is compatible with probabilities for members of a partition never settling down.

³¹ Jeffrey (1988) calls this “superconditioning”.

Theorem 7.4. is a variant of a similar result proved in Skyrms (1997). Skyrms' treatment of convergence is not confined to probability kinematics. It encompasses probabilistic black-box learning situations where an agent moves from a prior to a posterior via a learning process that remains unspecified. This is considerably more general than probability kinematics, which assumes that the evidential input is restricted to a partition. For the black-box setting, Skyrms shows that, while the limit need not be the indicator of the event, there is convergence to some limit, or to what Skyrms calls a *maximally informed opinion*.

§8. Principle (M) and Dynamic Coherence Skyrms' account of convergence is based on a martingale condition. This condition can be justified by a dynamic coherence argument. We will see in this section that condition (M) is a special case of that martingale condition. This will allow us to apply Skyrms' dynamic coherence argument to justify (M) and also to extend earlier results.

Skyrms' dynamic coherence argument is a generalizations of the dynamic Dutch book arguments given in Goldstein (1983) and van Fraassen (1984). It also applies to probability kinematics. Suppose that E is a member of \mathfrak{E}_n , and let $p_m = \mathbb{P}_m[E]$ for $m \geq n$. In general, moving from p_m to p_{m+1} can be thought of as a black-box learning experience. From your current point of view p_m could in principle take on any value between zero and one. We shall assume a bit more structure, however. Each p_m is a \mathfrak{F}_m -measurable random variable. This means that all propositions expressible by p_m are elements of \mathfrak{F}_m . The underlying probability space (Ω, \mathfrak{F}) is assumed to be sufficiently rich. Each element of Ω specifies a value for any p_m .

Dynamic coherence puts considerable constraints on a sequence of degrees of belief. Skyrms (1997) demonstrates that dynamic coherence requires a sequence such as p_n, p_{n+1}, \dots to be a martingale. There is one minor difference to our setting. In Skyrms' case, the sequence of σ -fields is generated by the beliefs to date, but a little reflection shows that this difference is inessential. The σ -algebra generated by the beliefs p_n, \dots, p_m is by assumption a sub- σ -algebra of \mathfrak{F}_m , so the propositions Skyrms requires to be expressible are certainly in \mathfrak{F}_m . Furthermore, since Ω is sufficiently rich, every proposition in \mathfrak{F}_m determines a range of possible values of p_m .

That dynamic coherence implies the martingale property means that

$$\int_G p_{m+1} d\mathbb{P} = \int_G p_m d\mathbb{P} \text{ for all } m \geq n \text{ and all } G \in \mathfrak{F}_m. \quad (\text{M}')$$

Thus (M') says that on average future degrees of belief are equal to present degrees of belief.³²

Suppose that the agent believes with probability one that each p_m will be equal to a constant k_m a.e. (\mathbb{P}). Then it follows from (M') (with $G = \Omega$) that

$$k_{m+1} = k_m.$$

³² Soft Jeffrey updates do not satisfy the martingale property (M') whenever the expectation of the soft update is not equal to the current marginal expectation. For an example see Seidenfeld (1987, Corollary 1) (this example is relevant because of the relationship between probability kinematics and minimum Kullback-Leibler shifts, which is also explained in that article).

Thus $p_{m+1} = p_m$ a.e. This is principle (M) with an almost surely qualification. What this means is that we can view (M) as a special case of (M') where an agent is certain that probability kinematics operates on constant probabilities for members of partitions.

As a more general example, suppose that the random variables p_n, p_{n+1}, \dots only take on values in the set $\{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$. Then (M') implies that

$$\begin{aligned} \frac{1}{4}\mathbb{P}\left[p_{m+1} = \frac{1}{4}\right] + \frac{1}{2}\mathbb{P}\left[p_{m+1} = \frac{1}{2}\right] + \frac{3}{4}\mathbb{P}\left[p_{m+1} = \frac{3}{4}\right] \\ = \frac{1}{4}\mathbb{P}\left[p_m = \frac{1}{4}\right] + \frac{1}{2}\mathbb{P}\left[p_m = \frac{1}{2}\right] + \frac{3}{4}\mathbb{P}\left[p_m = \frac{3}{4}\right] \end{aligned}$$

This is the sense in which probabilities for the members of a partition need to cohere over time. While they are allowed to change, they agree in prior expectation (even in the more stringent sense of (M') where the expectation can be taken with respect to any G in \mathfrak{F}_m).

The dynamic coherence arguments used by Goldstein, van Fraassen and Skyrms have been met with a lot of critical resistance in the literature (see e.g. Levi, 1987; Christensen, 1991; Talbott, 1991; Maher, 1992; Howson & Urbach, 1993). These criticisms are aimed at the so-called 'reflection principle', which is, essentially, an instance of (M'): it says that one's conditional probability for A given that one's anticipated future probability for A is r has to be equal to r as well.

It is easy to find counterexamples to this relation between current and future degrees of belief. Just think of Ulysses and the Sirens, taking drugs, forgetting what one knows, and so on. In all these cases something that need not be considered irrational leads to inconsistencies between current and anticipated future degrees of belief. However, concluding that the reflection principle is not a principle of rational belief change misses a point about reflection that is, more or less explicitly, inherent in the framework of Goldstein, van Fraassen and Skyrms.³³ Reflection is supposed to be a principle governing rational belief change *based on learning*. The counterexamples to reflection show that belief change may also have sources other than learning. But if an agent contemplates belief revision based on learning and nothing else, reflection seems to be a plausible principle. This view of the reflection principle is more fully worked out in Huttegger (2013, 2014), where alternative approaches to justifying reflection besides dynamic coherence are developed. One is based on minimizing expected accuracy, the other one on the idea that acquiring new beliefs from learning should not lead one to expecting to make bad decisions. Thus, there are good reasons besides dynamic coherence to view the sort of reflection principle expressed by (M') as directing rational belief change.

Before we are going to derive some consequences of (M'), one crucial underlying assumption should not go unmentioned. In an important paper Kadane et al. (1996) discuss the role of countable additivity for several issues surrounding the reflection principle. In particular, reflection can fail if a probability measure is only finitely additive.³⁴ The topic of finite versus countable additivity is a very involved

³³ See also Skyrms (1990) and van Fraassen (1995).

³⁴ More precisely, (M') is a condition of disintegrability of probabilities at time m with respect to the partition used to define probabilities at time $m + 1$. Dubins (1975) shows that with respect to bounded random variables disintegrability is equivalent to

one that I don't wish to enter here. At the very least, one can take countable additivity as a special assumption (which need not hold but often does hold as an approximation) for the results discussed earlier. Countable additivity is also an assumption underlying the Blackwell-Dubins theorem. So it is something that has to be assumed for the merging of opinions results considered here. For the convergence results (convergence to certainty, convergence to a maximally informed opinion), it appears that an alternative account in terms of certain finitely additive probability measures can be given. As Zabell (2002) points out, there is a martingale convergence theorem for finitely additive probability measures that are strategic (for details see Purves & Sudderth, 1976); this is also discussed in Seidenfeld (1985, Appendix). It seems possible that a finitely additive martingale theorem could also be applied to the cases considered by us.

To sum up, we see that principle (M) and its generalization (M') are far from being arbitrary. They turn out to follow from several quite natural coherence conditions. In the next section we show that some of our previous results find immediate extensions for (M').

§9. Generalizations In the previous section we have considered random probabilities p_n for members of \mathfrak{E}_n . If we have such a random probability for each member of \mathfrak{E}_n , this defines a random probability measure P_n on \mathfrak{F}_n . Let p_n^E be the random probability of $E \in \mathfrak{E}_n$ assigned by P_n , and let $p_m^E, m > n$ be the probabilities assigned to E by the measures $P_m, m > n$ on \mathfrak{F}_m . We say that the martingale condition (M') holds for a sequence of probability measures P_1, P_2, \dots on the measurable spaces $(\Omega, \mathfrak{F}_1), (\Omega, \mathfrak{F}_2), \dots$ if for all n and for all $E \in \mathfrak{E}_n$ (M') holds for p_n^E, p_{n+1}^E, \dots , that is

$$\int_G p_{m+1}^E d\mathbb{P} = \int_G p_m^E d\mathbb{P} \text{ for all } m \geq n \text{ and all } G \in \mathfrak{F}_m.$$

This setup describes a learning situation where an agent is uncertain about the probability distributions over partitions. This requires that the underlying probability space again be sufficiently rich, specifying values for each possible sequence of measures $P_n, n = 1, 2, \dots$.

Our first result generalizes convergence to a maximally informed opinion to this learning situation.

THEOREM 9.5. *Suppose that $P_n, n = 1, 2, \dots$ is a random sequence of probability measures on $(\Omega, \mathfrak{F}_1), (\Omega, \mathfrak{F}_2), \dots$, and that $P_n, n = 1, 2, \dots$ is uniformly absolutely continuous with respect to \mathbb{P} a.e. (\mathbb{P}) . Suppose also that (M') holds for $P_n, n = 1, 2, \dots$ a.e. (\mathbb{P}) . If $A \in \mathfrak{F}$, then $\mathbb{P}_n[A]$ converges to a random limit as $n \rightarrow \infty$ a.e. (\mathbb{P}) .*

The theorem presupposes that essentially all random probability measures on \mathfrak{F}_n that one contemplates are absolutely continuous relative to one's prior. In this case, (M') implies convergence to a maximally informed opinion.

conglomerability in a partition for bounded random variables. Countable additivity implies conglomerability in countable partitions. Countable probabilities may fail to be conglomerable in uncountably infinite partitions (Seidenfeld et al., 2014).

There also is an analogue to the merging result for hard Jeffrey shifts. P_n and Q_n are now viewed as random probability measures on \mathfrak{F}_n for each n , and Theorem 5.1. generalizes to:

THEOREM 9.6. *Suppose that $P_n, n = 1, 2, \dots$ and $Q_n, n = 1, 2, \dots$ are random sequences of probability measures on $(\Omega, \mathfrak{F}_1), (\Omega, \mathfrak{F}_2), \dots$, $Q_n = P_n$, that $Q_n, n = 1, 2, \dots$ is uniformly absolutely continuous with respect to \mathbb{Q} a.e. (\mathbb{Q}) , and that $\mathbb{Q} \ll \mathbb{P}$. If condition (M') holds for $Q_n, n = 1, 2, \dots$ a.e. (\mathbb{Q}) , then $d(\mathbb{P}_n, \mathbb{Q}_n) \rightarrow 0$ as $n \rightarrow \infty$ a.e. (\mathbb{Q}) .*

Under the assumptions of Theorem 9.6., \mathbb{P}_n merges to \mathbb{Q}_n in the sense that a.e. (\mathbb{Q}) the two measures get arbitrarily close in variational distance. Thus, by considering probability kinematics in terms of random probability distributions on partitions, we recover the Blackwell-Dubins notion of merging.

In §5. we illustrated Theorem 5.1. in terms of mechanical observers that provide both agents with joint probability distributions. Viewing the martingale condition (M') as in the previous section as a basic principle of rational learning is relevant for this interpretation. Recall that (M) was one of the assumptions of Theorem 5.1. (our first merging result). Now, (M) requires that the joint distributions of members of partitions don't change in response to later experiences. This is a severe restriction of learning situations; it requires that later evidence be evidentially independent of earlier outcomes. For mechanical observers, this does not appear to be a plausible assumption, at least not in general. However, (M) is only a special case of (M'). The more general martingale principle allows the kind of evidential dependence between earlier outcomes and later evidence that (M) prohibits. What (M') does constrain is the way agents think about this dependence; there is supposed to be no such dependence on average. As we put it in the previous section, agents think of the mechanical observer as a genuine source for learning. This is a normatively ideal case from which mechanical observers may deviate in many possible ways. One such class of scenarios was treated in Theorem 5.2., where outputs of the mechanical observer may violate the martingale condition yet still lead to merging. This result covers many, but by no means all, cases. One particularly important question, to which I don't yet have an answer, is whether the normative ideal can be approximated with mechanical observers that are not as trustworthy as the ideal one. The normatively ideal case is nevertheless important, for it represents a model of inquiry that obtains under the best epistemic conditions, one that we might be able to approximate to varying degrees.

The more general treatment outlined in this section may also be applied to the non-merging result for Bayes factors. The sequences of probability measures in Theorem 6.3. may be random, the only constraint being that the Bayes factor identities (6) hold almost surely. Since (M') requires that probabilities for elements of a partition, once adopted, will be the same on average in the future, it seems plausible that the non-merging result can be extended to the present setting. Proving this conjecture is beyond the reach of this essay, though.

§10. Concluding Remarks on Disagreement and Rationality According to Savage (1954), one of the main goals of statistics is to find ways for dealing with interpersonal differences in judgements. Statistics can bring diverging views closer to each other by the analysis of evidence. The merging of opinions theorem

due to Blackwell and Dubins is a very general expression of this idea within the framework of subjective Bayesian probability theory. Our first new insight is that merging of opinions generalizes to situations where evidence is uncertain but shared among agents in terms of a common joint probability distribution over experimental outcomes. These two merging results seem to be especially relevant for the alleged failure of subjective Bayesianism to account for the rationality of science. For a subjective Bayesian, the rationality of science consists in its ability to lead to a consensus under certain conditions. While our initial belief states can be quite arbitrary, successive rational belief revisions based on the same evidence lead us to have basically the same beliefs. Where we end up does not depend on where we begin as long as there is enough evidence.

It should be emphasized that consensus is expected *only under certain conditions*, and the claim is that those conditions often (or at least sometimes) do hold in scientific investigations. Our theorems suggest what those conditions are. They fall into two groups. In the first place, agents need to be *dynamically rational*. In the second place, their belief revisions need to operate in *appropriate circumstances*.

Dynamic rationality means that agents update by probability kinematics and obey principle (M) or (M'). The rationality of these requirements has to be understood in the right way. Justifications of probability kinematics presume that the partition used for probability kinematics is a sufficient partition; i.e., the learning situation is such that what is learned does not go beyond what is given by the partition. The martingale condition (M') further restricts the sequence of learning experiences to those that are diachronically rational from an agents prior point of view.

The dynamic rationality of agents is enough to guarantee convergence to certainty or convergence to a maximally informed opinion in the corresponding learning situation. But it is not enough to guarantee merging of opinions, which also requires suitable circumstances for learning. First, agents have to agree on probability zero events (absolute continuity of priors). Furthermore, posteriors cohere with priors in the sense given by uniform absolute continuity. Finally, agents change beliefs based on shared evidence. All these conditions are, implicitly or explicitly, invoked by our merging results. The Blackwell-Dubins result accounts for the case where evidence comes in terms of propositions. Our novel merging result deals with the case where the evidential input is given by a shared probability distribution over experimental outcomes. If an observational interaction has an appropriate structure, this can be modeled by probability kinematics with hard Jeffrey shifts.

In both cases the Bayesian strategy for explaining the rationality of science works. It has to be emphasized, though, that the conditions that are required for this to hold are far from being trivially satisfied. Our priors may fail to be absolutely continuous, and there may be no way you could persuade me to change my priors by appealing to my rationality. Similarly, we may disagree about the evidence we get, which could prompt our beliefs to diverge. So, even if we are individually rational, we may fail to agree in the long run unless these conditions are met.

Our second main insight concerns the circumstances of *non-merging*. In addition to the conditions that were just mentioned, both merging results assume rather solid kinds of evidential inputs. Although this might be a useful assumption for science, it is not a good one in general. Soft Jeffrey shifts are one option for describing more fluid kinds of evidential inputs for which we have a reasonable criterion as to when

two agents have the same input. In this case our beliefs can diverge in the long run even if all of the conditions mentioned before hold.

A related result was proved by Schervish & Seidenfeld (1990). Shervish & Seidenfeld don't assume that agents update by Bayesian conditioning; instead, posteriors are conditional probabilities chosen from a given set of probability measures. If this set has specific properties, then the agent's posteriors will merge with increasing information. However, this consequence depends on the properties of the set of probability measures. If the set is characterized by rather weak properties, then consensus is not guaranteed.³⁵ If we interpret the set of probability measures as consisting of those measures whose conditional probabilities are admissible responses to the same evidence, then this result says that if the set is too unconstrained, two agents responding to the same evidence may disagree forever.

Non-merging result are important for one aspect of the epistemological questions surrounding disagreement: the relationship between individual rationality and the possibility of consensus in the short and in the long run. There are other issues in addition to this. For instance, one might think about principles of *group rationality* that should govern the pooling of different opinions (e.g. Lehrer & Wagner, 1981). This goes beyond the concerns of this essay; one may well imagine pooling the opinions of individually rational agents who have the same evidential inputs. The principles governing the pooling go beyond the individual rationality of the agents. Seidenfeld et al. (1989) show that a complete Bayesian solution of the pooling problem is impossible. A more restricted treatment might be feasible, however (see Mongin, 1995).

Our results lead to the conclusion that, even under otherwise favorable circumstances, a soft kind of information allows individual rationality to be consistent with sustained disagreement. I don't think that this is a weakness of the broadly Bayesian approach advocated in this essay. Merging of beliefs happens when it should, i.e., under conditions which may, for example, hold for certain carefully designed scientific investigations. But the claim of merging is not a no-brainer that can be used across the board. In a well known essay, van Inwagen (1996) discusses three fields (philosophy, politics and religion—a list which could certainly be extended) where disagreements among well informed and qualified people abound. Given the background of the admittedly simple but useful model of soft Jeffrey shifts we should say that this kind of disagreement is to be expected. If uncertain evidence leaves room for interpretation, then individually rational people can reach different conclusions based on the same evidence. This appears to, at least partly, resolve what is puzzling about many situations where individuals acknowledge each others rationality in the face of disagreement but each one maintains their beliefs. From a Bayesian perspective, there is nothing wrong with this, except under the special conditions where merging results apply.

§11. Acknowledgements A version of this paper was presented at the University of Salzburg, where the audience provided many helpful comments. I wish

³⁵ For merging the set of probability measures has to be closed, convex, generated by finitely many extreme measures such that all measures in that set are mutually absolutely continuous. If the extreme points are only weak-star compact, then there is no assurance of consensus at all.

to thank Jim Joyce, Hannes Leitgeb, Samir Okasha, Jan-Willem Romeijn, Brian Skyrms, Bas van Fraassen, Carl Wagner, Kevin Zollman, an anonymous referee of this journal, and especially Teddy Seidenfeld, who prepared detailed commentaries.

§12. Appendix

12.1. Martingales, Convergence to Certainty, and Regular Conditional Distributions Throughout the probability measure \mathbb{P} is assumed to be a countably additive probability measure on a measurable space (Ω, \mathfrak{F}) . Let X be an \mathfrak{F} -measurable random variable. $E[X|\mathfrak{F}_n]$ is the conditional expectation of X given \mathfrak{F}_n . The expectation is taken with respect to \mathbb{P} . $E[X|\mathfrak{F}_n]$ is a \mathfrak{F}_n -measurable random variable and gives your estimate of the value of X if you know, for any event in \mathfrak{F}_n , whether or not it has occurred.

In our case each σ -field \mathfrak{F}_n is generated by a partition of Ω . If $\mathbb{P}[E] > 0$ for $E \in \mathfrak{E}_n$, it follows that

$$E[X|\mathfrak{F}_n](\omega) = \frac{1}{\mathbb{P}[E]} \int_E X d\mathbb{P} \quad (8)$$

for almost all ω in E (e.g. Ash, 2000, Section 5.5). This means that on each set of positive probability in the partition \mathfrak{E}_n , the conditional expectation $E[X|\mathfrak{F}_n]$ is almost surely constant.

A sequence of random variables Z_1, Z_2, \dots is a martingale with respect to $\mathfrak{F}_1, \mathfrak{F}_2, \dots$ if $E[Z_n|\mathfrak{F}_{n-1}] = Z_{n-1}$. That $E[X|\mathfrak{F}_n]$ is a martingale follows from a fundamental property of conditional expectations: $E[E[X|\mathfrak{F}_n]|\mathfrak{F}_{n-1}] = E[X|\mathfrak{F}_{n-1}]$. The martingale convergence theorem implies that almost surely

$$\lim_{n \rightarrow \infty} E[X|\mathfrak{F}_n] = E[X|\mathfrak{F}] = E[X|\mathfrak{F}] = X.$$

The last equality holds if X is \mathfrak{F} -measurable. It follows that the sequence of conditional probabilities $\mathbb{P}[A|\mathfrak{F}_1], \mathbb{P}[A|\mathfrak{F}_2], \dots$ converges to I_A whenever $A \in \mathfrak{F}$ (where I_A is the indicator function of A). Just let $X = I_A$. Then $E[X|\mathfrak{F}_n](\omega)$ is the conditional probability of A given \mathfrak{F}_n , i.e., $\mathbb{P}[A|\mathfrak{F}_n](\omega)$.

If each member of \mathfrak{E}_n has positive probability, for all n , then by equation (8),

$$\mathbb{P}[A|\mathfrak{F}_n](\omega) = \frac{\mathbb{P}[A \cap E_n(\omega)]}{\mathbb{P}[E_n(\omega)]} = \mathbb{P}[A|E_n(\omega)] \quad (9)$$

a.e.. The martingale convergence theorem then implies that the sequence of conditional probabilities $\mathbb{P}[A|E_n(\omega)]$ goes to zero or one almost surely.

Let \mathfrak{G} be a sub- σ -field of a σ -field \mathfrak{F} . The conditional probability $\mathbb{P}[A|\mathfrak{G}]$ of an event A given \mathfrak{G} exists for any event A and is almost surely unique. However, this does not imply that versions of each $\mathbb{P}[A|\mathfrak{G}]$ can be so chosen so that $\mathbb{P}[\cdot|\mathfrak{G}]$ is a countably additive probability measure over \mathfrak{G} for each fixed ω (e.g. Ash, 2000). There is a counterexample where the set for which $\mathbb{P}[\cdot|\mathfrak{G}]$ is not a countably additive probability measure has positive probability whatever version of $\mathbb{P}[A|\mathfrak{G}]$ is chosen (see Seidenfeld, 2001, and references therein).

If it is possible to choose versions of each $\mathbb{P}[A|\mathfrak{G}]$ so that $\mathbb{P}[\cdot|\mathfrak{G}]$ is a countably additive probability measure over \mathfrak{F} for every ω then $\mathbb{P}[\cdot|\mathfrak{G}]$ is a *regular conditional distribution*. Regular conditional distributions play an important role in the Blackwell-Dubins theorem, which assumes that the prior probability is *predictive*;

this means that the prior is assumed to admit a regular conditional distribution of the future given each finite history of observations.

Moreover, as shown in Blackwell & Dubins (1975) and Seidenfeld (2001) there can be a tension between regular conditional distributions and coherence. Even if a regular conditional distribution exists it need not have the conditioning event as its support; i.e. it might not be *proper*. (A regular conditional distribution \mathbb{P} given \mathfrak{G} is proper at ω if $\mathbb{P}[A|\mathfrak{G}](\omega) = 1$ whenever $\omega \in A \in \mathfrak{G}$, and it is proper if it is proper at each ω .)

Regular conditional distributions and properness lead to further qualifications for the Blackwell-Dubins theorem. A well-known sufficient condition for the existence of a regular conditional probability is that Ω is a separable completely metrizable topological space and \mathfrak{F} is its Borel- σ -algebra. Concerning properness, a regular conditional distribution is proper with probability one whenever the sub- σ -field \mathfrak{G} is countably generated (Seidenfeld, 2001, Theorem 1).

These conditions for the merging of opinions theorem are more general than what is used in the main text. In particular, they allow the application of continuous random variables (under the appropriate restrictions).

12.2. Proof of Theorem 5.1. For each $A \in \mathfrak{F}$, consider the random variable $\mathbb{P}[A|\mathfrak{F}_n]$. Probability kinematics calculates the new probability for A by taking the expected value of $\mathbb{P}[A|\mathfrak{F}_n]$ with respect to P_n :

$$\mathbb{P}_n[A] = \int_{\Omega} \mathbb{P}[A|\mathfrak{F}_n](\omega) P_n(d\omega).$$

This is equivalent to (4). To see this, note that $\mathbb{P}[A|\mathfrak{F}_n] = \sum_{E \in \mathfrak{E}_n} I_E \mathbb{P}[A|E]$ a.e. (\mathbb{P}). Let S be the set of all ω for which this relation holds. Then $\mathbb{P}[S] = 1$. Because $S \in \mathfrak{F}_n$, we also have $P_n[S] = 1$ by absolute continuity. Thus the integral above is essentially an integral of a simple function which is equal to (4).

If $Q_n = P_n$ on \mathfrak{F}_n for each n , then $dQ_n = dP_n$; so

$$|\mathbb{P}_n[A] - Q_n[A]| = \left| \int_{\Omega} \mathbb{P}[A|\mathfrak{F}_n] dP_n - \int_{\Omega} Q[A|\mathfrak{F}_n] dQ_n \right| \leq \int_{\Omega} |\mathbb{P}[A|\mathfrak{F}_n] - Q[A|\mathfrak{F}_n]| dQ_n.$$

Letting $\mathbb{P}_{\mathfrak{F}_n}(\cdot) = \mathbb{P}[\cdot|\mathfrak{F}_n]$, $Q_{\mathfrak{F}_n}(\cdot) = Q[\cdot|\mathfrak{F}_n]$ and taking supremums on both sides gives

$$d(\mathbb{P}_n, Q_n) = \sup_A |\mathbb{P}_n[A] - Q_n[A]| \leq \sup_A \int_{\Omega} |\mathbb{P}[A|\mathfrak{F}_n] - Q[A|\mathfrak{F}_n]| dQ_n. \quad (10)$$

Clearly, $|\mathbb{P}[A|\mathfrak{F}_n] - Q[A|\mathfrak{F}_n]| \leq \sup_A |\mathbb{P}[A|\mathfrak{F}_n] - Q[A|\mathfrak{F}_n]|$ for all $\omega \in \Omega$. Hence,

$$\int_{\Omega} |\mathbb{P}[A|\mathfrak{F}_n] - Q[A|\mathfrak{F}_n]| dQ_n \leq \int_{\Omega} \sup_A |\mathbb{P}[A|\mathfrak{F}_n] - Q[A|\mathfrak{F}_n]| dQ_n.$$

Since this inequality holds for each A , we have

$$\sup_A \int_{\Omega} |\mathbb{P}[A|\mathfrak{F}_n] - Q[A|\mathfrak{F}_n]| dQ_n \leq \int_{\Omega} \sup_A |\mathbb{P}[A|\mathfrak{F}_n] - Q[A|\mathfrak{F}_n]| dQ_n.$$

It now follows from (10) that

$$d(\mathbb{P}_n, Q_n) \leq \int_{\Omega} d(\mathbb{P}_{\mathfrak{F}_n}, Q_{\mathfrak{F}_n}) dQ_n. \quad (11)$$

That $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})$ is a \mathfrak{F}_n -measurable random variable follows from the integral representation of d given in Blackwell & Dubins (1962, p. 883).

Since $Q_n \ll \mathbb{Q}$ on the measurable space (Ω, \mathfrak{F}_n) , the Radon-Nikodým derivative $Y_n = dQ_n/d\mathbb{Q}$ exists. Each Y_n is \mathfrak{F}_n -measurable and

$$Q_n(F) = \int_F Y_n d\mathbb{Q} \quad \text{for each } F \in \mathfrak{F}_n.$$

Moreover, the integral on the right side of (11) is equal to $\int_{\Omega} d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) Y_n d\mathbb{Q}$ (the same is true for every $F \in \mathfrak{F}_n$), so that

$$d(\mathbb{P}_n, \mathbb{Q}_n) \leq \int_{\Omega} d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) Y_n d\mathbb{Q}. \quad (12)$$

Our goal is to show that the right side of this equation goes to zero.

The next lemmata note some basic facts about the sequence Y_1, Y_2, \dots of Radon-Nikodým derivatives. For the first, recall that a family of nonnegative random variables is uniformly integrable if for each $\varepsilon > 0$ there exists a K such that

$$\int_{\{Z \geq K\}} Z d\mathbb{Q} < \varepsilon,$$

for any Z in the family of random variables (K does not depend on Z).

LEMMA 12.7. *Let the sequence $Q_n, n = 1, 2, \dots$ be uniformly absolutely continuous with respect to \mathbb{Q} . Then $Y_n, n = 1, 2, \dots$ is uniformly integrable.*

Proof. Note that for any $K > 0$

$$\int_{\{Y_n \geq K\}} Y_n d\mathbb{Q} = \int_{\{Y_n \geq K\}} dQ_n = Q_n[Y_n \geq K]. \quad (13)$$

(Since Y_n is \mathfrak{F}_n -measurable, $\{Y_n \geq K\}$ is clearly an event in \mathfrak{F}_n .) By Markov's inequality, we have

$$\mathbb{Q}[Y_n \geq K] \leq \frac{E[Y_n]}{K} = \frac{1}{K}. \quad (14)$$

Since the sequence $Q_n, n = 1, 2, \dots$ is assumed to be a uniformly absolutely continuous family with respect to \mathbb{Q} , we know that for every $\varepsilon > 0$ there exists a $\delta > 0$ such that for all n

$$\mathbb{Q}[B] < \delta \implies Q_n(B) < \varepsilon \quad (15)$$

for all $B \in \mathfrak{F}_n$. Now, let $\varepsilon > 0$ and choose a $\delta > 0$ such that (15) holds uniformly in n and $B \in \mathfrak{F}_n$. If we choose K sufficiently large so that $1/K < \delta$ (this does not depend on n), then by (14) for all n

$$\mathbb{Q}[Y_n \geq K] < \delta.$$

Thus, by (15), $Q_n[Y_n \geq K] < \varepsilon$ for all n , and it follows from (13) that

$$\int_{\{Y_n \geq K\}} Y_n d\mathbb{Q} < \varepsilon$$

uniformly in n . Hence Y_1, Y_2, \dots is uniformly integrable. \square

LEMMA 12.8. *Let the sequence $Y_n, n = 1, 2, \dots$ be uniformly absolutely continuous with respect to \mathbb{Q} . Then the sequence $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) Y_n, n = 1, 2, \dots$ is uniformly integrable.*

Proof. By Lemma 12.7., the sequence $Y_n, n = 1, 2, \dots$ is uniformly integrable. It is well known that a family of random variables that is dominated by a uniformly integrable family of random variables is itself uniformly integrable. The sequence $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})Y_n, n = 1, 2, \dots$ is dominated by $Y_n, n = 1, 2, \dots$ since $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})Y_n \leq Y_n$ a.e. for all n ($d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) \leq 1$). Hence $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})Y_n, n = 1, 2, \dots$ is uniformly integrable. \square

LEMMA 12.9. *If condition (M) holds, then Y_1, Y_2, \dots is a martingale with respect to the sequence of fields $\mathfrak{F}_1, \mathfrak{F}_2, \dots$*

Proof. Suppose that $F \in \mathfrak{F}_{n-1}$. Then

$$\int_F Y_{n-1} d\mathbb{Q} = Q_{n-1}[F].$$

Since also $F \in \mathfrak{F}_n$

$$\int_F Y_n d\mathbb{Q} = Q_n[F].$$

If (M) holds, then

$$\int_F Y_n d\mathbb{Q} = \int_F Y_{n-1} d\mathbb{Q}$$

for all $F \in \mathfrak{F}_{n-1}$. It follows that Y_{n-1} is a version of the conditional expectation $E[Y_n | \mathfrak{F}_{n-1}]$. Because this is true for all n , Y_1, Y_2, \dots is a martingale. \square

COROLLARY. The sequence Y_1, Y_2, \dots has a finite limit a.e. (\mathbb{Q}).

Proof. This follows from the martingale convergence theorem since all Y_n are non-negative. \square

To finish the proof, observe that $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) \rightarrow 0$ a.e. (\mathbb{Q}) by Blackwell-Dubins since $\mathbb{Q} \ll \mathbb{P}$. Hence, by the corollary of Lemma 12.9., the sequence

$$d(\mathbb{P}_{\mathfrak{F}_1}, \mathbb{Q}_{\mathfrak{F}_1})Y_1, d(\mathbb{P}_{\mathfrak{F}_2}, \mathbb{Q}_{\mathfrak{F}_2})Y_2, \dots$$

converges to 0 a.e. (\mathbb{Q}). Together with Lemma 12.7. and Lemma 12.8., the Vitali convergence theorem (e.g. Ash (2000, Theorem 6.5.2 b)) now implies that

$$\lim_{n \rightarrow \infty} \int_{\Omega} d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})Y_n d\mathbb{Q} = \int_{\Omega} \lim_{n \rightarrow \infty} d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})Y_n d\mathbb{Q} = 0. \quad (16)$$

From this and (12) it follows that $d(\mathbb{P}_n, \mathbb{Q}_n) \rightarrow 0$ as $n \rightarrow \infty$.

12.3. Proof of Theorem 5.2. $\mathbb{Q}_n \ll \mathbb{Q}|_{\mathfrak{F}_n}$ implies that Y_n exists. We have $Y_n \leq g$ for some non-negative function g that is integrable:

$$\int_{\Omega} g d\mathbb{Q} < \infty$$

Hence, $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})Y_n \leq d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})g$. Also, since $\mathbb{Q} \ll \mathbb{P}$, $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) \rightarrow 0$ a.e. (\mathbb{Q}). Thus, because g is finite a.e. (otherwise its integral would diverge), we must have $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})g \rightarrow 0$ a.e. Therefore $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})Y_n \rightarrow 0$ a.e. Equation (16) now follows from the dominated convergence theorem because $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})Y_n \leq g$, and the conclusion of the theorem follows again from combining this with (12).

12.4. Proof of Theorem 6.3. We start with noting that Wagner (2002, p. 271) shows that the following lemma is true:

LEMMA 12.10. *Condition (6) holds if and only if there exists a c_n such that for all $E \in \mathfrak{E}_n$:*

$$\mathfrak{R}_{\mathbb{P}_n, \mathbb{P}_{n-1}}(E) = c_n \mathfrak{R}_{\mathbb{Q}_n, \mathbb{Q}_{n-1}}(E) \quad (17)$$

Proof. Suppose that such a c_n exists. Then (6) holds since

$$\mathfrak{B}_{\mathbb{P}_n, \mathbb{P}_{n-1}}(E, F) = \frac{\mathfrak{R}_{\mathbb{P}_n, \mathbb{P}_{n-1}}(E)}{\mathfrak{R}_{\mathbb{P}_n, \mathbb{P}_{n-1}}(F)} = \frac{\mathfrak{R}_{\mathbb{Q}_n, \mathbb{Q}_{n-1}}(E)}{\mathfrak{R}_{\mathbb{Q}_n, \mathbb{Q}_{n-1}}(F)} = \mathfrak{B}_{\mathbb{Q}_n, \mathbb{Q}_{n-1}}(E, F).$$

Suppose now that (6) holds, and define c_n by

$$c_n = \frac{\mathbb{P}_n[E_1]}{\mathbb{P}_{n-1}[E_1]} \bigg/ \frac{\mathbb{Q}_n[E_1]}{\mathbb{Q}_{n-1}[E_1]},$$

where $E_1 \in \mathfrak{E}_n$. Then it follows from (6) that (17) holds for all $E \in \mathfrak{E}_n$. \square

The first part of Theorem 6.3. asserts that condition (17) is consistent with condition (M) applied to both sequences of probability measures. This can be demonstrated by showing that certain systems of linear equations have solutions. Let $n \geq 2$ be arbitrary and F be in \mathfrak{E}_n . Then, by assumption (N1), there are events $E_1, E_2 \in \mathfrak{E}_{n+1}$ such that $F = E_1 \cup E_2$. Let $\mathbb{P}_n[F] = s$ and $\mathbb{Q}_n[F] = t$. Also, let $\mathbb{P}_n[E_1] = p$, $\mathbb{P}_n[E_2] = s - p$, $\mathbb{P}_{n+1}[E_1] = x_1$ and $\mathbb{P}_{n+1}[E_2] = x_2$. Similarly, let $\mathbb{Q}_n[E_1] = q$, $\mathbb{Q}_n[E_2] = t - q$, $\mathbb{Q}_{n+1}[E_1] = y_1$ and $\mathbb{Q}_{n+1}[E_2] = y_2$.

Both (17) and (M) are satisfied concerning F, E_1, E_2 if the following system of linear equations in the variables x_1, x_2, y_1, y_2 has nonnegative solutions:

$$\begin{aligned} \frac{1}{p}x_1 - \frac{c}{q}y_1 &= 0 \\ \frac{1}{s-p}x_2 - \frac{c}{t-q}y_2 &= 0 \\ x_1 + x_2 &= s \\ y_1 + y_2 &= t \end{aligned}$$

The solution for this system is given by

$$\begin{aligned} x_1 &= \frac{p(s(q-t) + ct(s-p))}{qs - pt}, \\ y_1 &= \frac{q(s(q-t) + ct(s-p))}{c(qs - pt)}, \end{aligned}$$

as well as $x_2 = s - x_1$ and $y_2 = t - y_1$. These solutions can be positive if $p/q \neq s/t$. This inequality follows from assumption (N3):

$$\frac{s}{t} = \frac{\mathbb{P}_n[F]}{\mathbb{Q}_n[F]} \neq \frac{\mathbb{P}[E_1|F]\mathbb{P}_n[F]}{\mathbb{Q}[E_1|F]\mathbb{Q}_n[F]} = \frac{\mathbb{P}_n[E_1]}{\mathbb{Q}_n[E_1]} = \frac{p}{q}.$$

If $p/q > s/t$, then x_1, y_1 are positive and strictly less than s, t , respectively, if

$$\frac{qs}{pt} < c < \frac{s(t-q)}{t(s-p)}. \quad (18)$$

If $p/q < s/t$, these inequalities need to be reversed.

What remains to be shown is that the same c can be chosen for any $F \in \mathfrak{E}_n$ (of which there are 2^n). To see that this is possible observe first that solutions can be calculated as above for any element of \mathfrak{E}_n . Moreover, observe that (18) defines an open interval including $c = 1$. We have

$$\frac{qs}{pt} = \frac{\mathbb{Q}_n[E_1]\mathbb{P}_n[F]}{\mathbb{P}_n[E_1]\mathbb{Q}_n[F]} = \frac{\mathbb{Q}[E_1|F]\mathbb{Q}_n[F]\mathbb{P}_n[F]}{\mathbb{P}[E_1|F]\mathbb{P}_n[F]\mathbb{Q}_n[F]} = \frac{\mathbb{Q}[E_1|F]}{\mathbb{P}[E_1|F]}$$

and similarly

$$\frac{s(t-q)}{t(s-p)} = \frac{\mathbb{Q}[E_2|F]}{\mathbb{P}[E_2|F]}.$$

By assumption (N3) $\mathbb{Q}[E_1|F]/\mathbb{P}[E_1|F] \neq 1$. If $\mathbb{Q}[E_1|F]/\mathbb{P}[E_1|F] > 1$, then $\mathbb{Q}[E_2|F]/\mathbb{P}[E_2|F] < 1$, and if $\mathbb{Q}[E_1|F]/\mathbb{P}[E_1|F] < 1$, then $\mathbb{Q}[E_2|F]/\mathbb{P}[E_2|F] > 1$. Hence, for each $F \in \mathfrak{E}_n$ there is an open interval around 1 that includes the admissible c values. Since \mathfrak{E}_n has only finitely many members, by taking the maximum lower value and the minimum higher value we get an open interval of c values that work for any F .

Consider now the partition $\mathfrak{E}_1 = \{E, \bar{E}\}$, and let $\mathbb{P}[E] = p$, $\mathbb{Q}[E] = q$, $\mathbb{P}_1[E] = p'$, $\mathbb{Q}_1[E] = q'$. Then by the same argument as above, the following system of linear equations must hold:

$$\begin{aligned} \frac{p'}{p} &= c \frac{q'}{q} \\ \frac{1-p'}{1-p} &= c \frac{1-q'}{1-q} \end{aligned}$$

By assumption (N2), $0 < p, q < 1$ and $p \neq q$. There exist positive solutions for p' and q' such that p' and q' are strictly between zero and one provided that

$$\frac{q}{p} < c < \frac{1-q}{1-p}.$$

whenever $p > q$ (otherwise the inequalities have to be reversed). This, together with the above argument, proves the assertion that there are sequences P_1, P_2, \dots and Q_1, Q_2, \dots that satisfy (6) and (M).

The second assertion of the theorem now follows immediately. Let $\varepsilon = |p' - q'|$. Then c can be chosen such that $\varepsilon > 0$. Because of condition (M), $\mathbb{P}_n[E] = \mathbb{P}_1[E] = p'$ and $\mathbb{Q}_n[E] = \mathbb{Q}_1[E] = q'$ for all $n = 1, 2, \dots$. Hence, $|\mathbb{P}_n[E] - \mathbb{Q}_n[E]| = \varepsilon, n = 1, 2, \dots$ and $d(\mathbb{P}_n, \mathbb{Q}_n) \geq \varepsilon > 0$ for all $n \geq 1$.

12.5. Proof of Theorem 7.4. The theorem assumes that the sequence $P_n, n = 1, 2, \dots$ is uniformly absolutely continuous with respect to \mathbb{P} . Therefore, $P_n \ll \mathbb{P}$ over each space (Ω, \mathfrak{F}_n) , and the Radon-Nikodým derivative $X_n = dP_n/d\mathbb{P}$ exists.

Suppose that condition (M) holds for the sequence of probability measures P_1, P_2, \dots on $(\Omega, \mathfrak{F}_1), (\Omega, \mathfrak{F}_2), \dots$. Then, by Lemma 12.9. and its corollary, X_1, X_2, \dots is a non-negative martingale that converges to a finite limit X a.e. Furthermore, as explained in 12.1., the sequence $\mathbb{P}[A|\mathfrak{F}_1], \mathbb{P}[A|\mathfrak{F}_2], \dots$ also is a martingale that converges to the indicator I_A a.e. It follows that $\mathbb{P}[A|\mathfrak{F}_n]X_n$ converges to $I_A X$ a.e.

Lemma 12.7. implies that X_1, X_2, \dots is uniformly integrable. A proof similar to the one of Lemma 12.8. shows that $\mathbb{P}[A|\mathfrak{F}_1]X_1, \mathbb{P}[A|\mathfrak{F}_2]X_2, \dots$ is uniformly integrable (just observe that $\mathbb{P}[A|\mathfrak{F}_n]X_n \leq X_n$). Now the Vitali convergence theorem is

applicable and we have

$$\lim_{n \rightarrow \infty} \int_{\Omega} \mathbb{P}[A|\mathfrak{F}_n] X_n d\mathbb{P} = \int_{\Omega} \lim_{n \rightarrow \infty} \mathbb{P}[A|\mathfrak{F}_n] X_n d\mathbb{P} = \int_A X d\mathbb{P}.$$

Since $\mathbb{P}_n[A] = \int_{\Omega} \mathbb{P}[A|\mathfrak{F}_n] X_n d\mathbb{P}$, it follows that $\mathbb{P}_n[A]$ converges to a limit.

12.6. Proof of Theorem 9.5. For any $A \in \mathfrak{F}$, let

$$\mathbb{P}_n[A](\omega) = \int_{\Omega} \mathbb{P}[A|\mathfrak{F}_n] dP_n(\omega),$$

This is equivalent to (4). Since $P_n, n = 1, 2, \dots$ is uniformly absolutely continuous relative to \mathbb{P} a.e. (\mathbb{P}) , the Radon-Nikodým derivatives $X_n = dP_n/d\mathbb{P}$ exist almost surely and it follows that

$$\mathbb{P}_n[A](\omega_1) = \int_{\Omega} \mathbb{P}[A|\mathfrak{F}_n](\omega_2) X_n(\omega_1, \omega_2) \mathbb{P}(d\omega_2) \quad (19)$$

for almost every ω_1 .

LEMMA 12.11. *If condition (M') holds for $P_n, n = 1, 2, \dots$ a.e. (\mathbb{P}) , then $X_n(\omega)$ converges to a finite limit a.e. (\mathbb{P}) for almost every ω .*

Proof. We can choose $\mathbb{P}[A|\mathfrak{F}_n]$ to be a regular conditional probability. Since \mathfrak{F}_n is generated by a finite partition, $\mathbb{P}[A|\mathfrak{F}_n]$ is also proper. Hence, if $F \in \mathfrak{F}_n$, then almost surely $\mathbb{P}[F|\mathfrak{F}_n](\omega) = 1$ if $\omega \in F$ and $\mathbb{P}[F|\mathfrak{F}_n](\omega) = 0$ if $\omega \notin F$. Hence, $\mathbb{P}[F|\mathfrak{F}_n] = I_F$ a.e., and so we get from (19) that

$$\mathbb{P}_n[F](\omega_1) = \int_F X_n(\omega_1, \omega_2) \mathbb{P}(d\omega_2)$$

for almost every ω_1 . Since (M') holds for $P_n, n = 1, 2, \dots$ a.e., we have for every $G, F \in \mathfrak{F}_m$

$$\int_G \int_F X_{m+1}(\omega_1, \omega_2) \mathbb{P}(d\omega_2) \mathbb{P}(d\omega_1) = \int_G \int_F X_m(\omega_1, \omega_2) \mathbb{P}(d\omega_2) \mathbb{P}(d\omega_1).$$

(Use the fact that F is a finite union of members of \mathfrak{E}_m .) Fubini's theorem implies that the order of integration can be reversed. Let \mathbb{P}^2 denote the product measure $\mathbb{P} \times \mathbb{P}$, and let H be any set in $\mathfrak{F}_m \times \mathfrak{F}_m$. Then

$$\int_H X_{m+1}(\omega_1, \omega_2) \mathbb{P}^2(d\omega_1, d\omega_2) = \int_H X_m(\omega_1, \omega_2) \mathbb{P}^2(d\omega_1, d\omega_2),$$

(That H does not need to be of the form $F \times G$ for $F, G \in \mathfrak{F}_m$ follows from a standard application of the monotone class theorem (Ash, 2000, Theorem 1.3.9).) Therefore, the sequence X_1, X_2, \dots is a martingale that converges to a finite limit a.e. relative to the product measure \mathbb{P}^2 .

It follows from this that, for almost every ω_1 , $X_m(\omega_1, \omega_2)$ converges to a finite limit for almost every ω_2 . For otherwise there would exist a set N_1 of positive measure $\mathbb{P}[N_1] > 0$ such that for every $\omega_1 \in N_1$ $X_m(\omega_1, \omega_2)$ does not converge to a finite limit for every ω_2 in a set N_2 also of positive measure $\mathbb{P}[N_2] > 0$. Since $\mathbb{P}^2[N_1 \times N_2] = \mathbb{P}[N_1]\mathbb{P}[N_2] > 0$, it would follow that there is a set $N_1 \times N_2$ of positive measure \mathbb{P}^2 where X_m does not converge to a finite limit, which contradicts the foregoing result. \square

Lemma 12.11. implies that the sequence of random variables

$$\mathbb{P}[A|\mathfrak{F}_n]X_n(\omega), n \geq 1$$

converges, for almost every ω , to a finite limit $I_A X(\omega)$ a.e. (\mathbb{P}).

What remains to be done in order to apply the Vitali convergence theorem is to show that the sequence $\mathbb{P}[A|\mathfrak{F}_n]X_n(\omega)$ is uniformly integrable for almost every ω . Since $P_n, n = 1, 2, \dots$ is uniformly absolutely continuous relative to \mathbb{P} a.e., a proof analogous to that of Lemma 12.7. shows that $X_n(\omega), n = 1, 2, \dots$ is uniformly integrable for almost every ω . Similarly, because $\mathbb{P}[A|\mathfrak{F}_n] \leq 1$, $\mathbb{P}[A|\mathfrak{F}_n]X_n(\omega), n = 1, 2, \dots$ is uniformly integrable for almost every ω by an argument along the lines of Lemma 12.8.

Putting things together, we can apply the Vitali convergence theorem in order to conclude that

$$\lim_{n \rightarrow \infty} \mathbb{P}_n[A](\omega) = \lim_{n \rightarrow \infty} \int_{\Omega} \mathbb{P}[A|\mathfrak{F}_n]X_n(\omega) d\mathbb{P} = \int_A X(\omega) d\mathbb{P}$$

for almost every ω .

12.7. Proof of Theorem 9.6. The proof is a combination of the proofs of Theorem 5.1. and Theorem 9.5. For $A \in \mathfrak{F}$ define

$$\mathbb{Q}_n[A](\omega) = \int_{\Omega} \mathbb{Q}[A|\mathfrak{F}_n] dQ_n(\omega).$$

Since $Q_n \ll \mathbb{Q}$ on (Ω, \mathfrak{F}) a.e. (\mathbb{Q}), the Radon-Nikodým derivative Y_n exists almost surely for each n . Lemma 12.11. implies that $Y_n(\omega), n = 1, 2, \dots$ converges to a finite limit almost surely.

The assumption that $Q_n = P_n$ a.e. (\mathbb{Q}) implies that

$$\begin{aligned} |\mathbb{P}_n[A](\omega) - \mathbb{Q}_n[A](\omega)| &= \left| \int_{\Omega} \mathbb{P}[A|\mathfrak{F}_n] dP_n(\omega) - \int_{\Omega} \mathbb{Q}[A|\mathfrak{F}_n] dQ_n(\omega) \right| \\ &\leq \int_{\Omega} |\mathbb{P}[A|\mathfrak{F}_n] - \mathbb{Q}[A|\mathfrak{F}_n]| dQ_n(\omega) \end{aligned}$$

for almost every ω (\mathbb{Q}). Taking supremums on both sides yields

$$d(\mathbb{P}_n(\omega), \mathbb{Q}_n(\omega)) \leq \int_{\Omega} d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) dQ_n(\omega) = \int_{\Omega} d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) Y_n(\omega) d\mathbb{Q}.$$

for almost every ω (\mathbb{Q}). (The equality follows from the fact that $Y_n(\omega)$ is the Radon-Nikodým derivative of $dQ_n(\omega)/d\mathbb{Q}$.)

The proof can now be finished along the lines of the proof of Theorem 5.1. by adding almost surely qualifications at the appropriate places. Proofs analogous to the ones of Lemma 12.7. and Lemma 12.8. show that $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})Y_n(\omega), n = 1, 2, \dots$ is a uniformly integrable sequence for almost every ω . Because $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) \rightarrow 0$ as $n \rightarrow \infty$ a.e. (\mathbb{Q}), $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n})Y_n(\omega) \rightarrow 0$ a.e. (\mathbb{Q}) for almost every ω . Hence, for almost every ω ,

$$\lim_{n \rightarrow \infty} \int_{\Omega} d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) Y_n(\omega) d\mathbb{Q} = \int_{\Omega} \lim_{n \rightarrow \infty} d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) Y_n(\omega) d\mathbb{Q} = 0.$$

It follows that $\lim_{n \rightarrow \infty} d(\mathbb{P}_n(\omega), \mathbb{Q}_n(\omega)) = 0$ for almost every ω (\mathbb{Q}).

- Armendt, B. (1980). Is there a Dutch book argument for probability kinematics? *Philosophy of Science* **47**, 583–588.
- Ash, R. B. (2000). *Probability and Measure Theory*. San Diego: Academic Press.
- Aumann, R. J. (1976). Agreeing to disagree. *The Annals of Statistics* **4**, 1236–1239.
- Billingsley, P. (2008). *Probability and Measure*. John Wiley & Sons.
- Blackwell, D., & Dubins, L. (1962). Merging of opinions with increasing information. *The Annals of Mathematical Statistics* **33**, 882–886.
- Blackwell, D., & Dubins, L. (1975). On existence and non-existence of proper, regular, conditional distributions. *Annals of Probability* **3**, 741–752.
- Blackwell, D., & Girshick, M. A. (1954). *Theory of Games and Statistical Decisions*. John Wiley & Sons.
- Bradley, R. (2005). Radical probabilism and Bayesian conditioning. *Philosophy of Science* **72**, 342–364.
- Carnap, R. (1950). *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Carnap, R. (1971). A basic system of inductive logic, part 1. In Rudolf Carnap & Richard C. Jeffrey, editors, *Studies in Inductive Logic and Probability I*, pp. 33–165. Los Angeles: University of California Press.
- Carnap, R. (1980). A basic system of inductive logic, part 2. In Jeffrey, R. C., editor, *Studies in Inductive Logic and Probability II*, pp. 7–155. Los Angeles: University of California Press.
- Chalmers, A. F. (1999). *What is This Thing Called Science*. Indianapolis: Hackett.
- Christensen, D. (1991). Clever bookies and coherent beliefs. *Philosophical Review* **100**, 229–247.
- Christensen, D. (2007). Epistemology of disagreement: the good news. *Philosophical Review* **116**, 187–217.
- Christensen, D. (2009). Disagreement as evidence: The epistemology of controversy. *Philosophy Compass* **4/5**, 756–767.
- D’Aristotile, A., Diaconis, P., & Freedman, D. (1988). On merging of probabilities. *Sankhiā: The Indian Journal of Statistics* **50**, 363–380.
- de Finetti, B. (1937). La prevision: ses lois logiques ses sources subjectives. *Annales d’Institut Henri Poincaré* **7**, 1–68. Translated in Kyburg, H. E. and Smokler, H. E., editors, *Studies in Subjective Probability*, pages 93–158, Wiley, New York, 1964.
- Dewey, J. (1929). *The Quest for Certainty*. Balch New York: Minton.
- Diaconis, P., & Freedman, D. (1986). On the consistency of Bayes estimates. *Annals of Statistics* **14**, 1–26.
- Diaconis, P., & Zabell, S. L. (1982). Updating subjective probability. *Journal of the American Statistical Association* **77**, 822–830.
- Döring, F. (1999). Why Bayesian psychology is incomplete. *Philosophy of Science* **66**, 379–389.
- Dubins, L. (1975). Finitely additive conditional probabilities, conglomerability, and disintegration. *Annals of Probability* **3**, 89–99.
- Earman, J. (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.
- Easwaran, K. (2013). Expected accuracy supports conditionalization—and conglomerability and reflection. *Philosophy of Science* **80**, 119–142.
- Elga, A. (2007). Reflection and disagreement. *Noûs* **41**, 478–502.

- Field, H. (1978). A note on Jeffrey conditionalization. *Philosophy of Science* **45**, 361–367.
- Gaifman, H., & Snir, M. (1982). Probabilities over rich languages, testing and randomness. *The Journal of Symbolic Logic* **47**, 495–548.
- Goldstein, M. (1983). The prevision of a prevision. *Journal of the American Statistical Association* **78**, 817–819.
- Good, I. J. (1950). *Probability and the Weighing of Evidence*. London: Charles Griffin.
- Good, I. J. (1983). *Good Thinking. The Foundations of Probability and Its Applications*. Minneapolis: University of Minnesota Press.
- Greaves, H., & Wallace, D. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind* **115**, 607–632.
- Gutting, G. (1982). *Religious Belief and Religious Scepticism*. Notre Dame: University of Notre Dame Press.
- Howson, C., & Urbach, P. (1993). *Scientific Reasoning. The Bayesian Approach* (second ed.). La Salle, Illinois: Open Court.
- Huttegger, S. M. (2013). In defense of reflection. *Philosophy of Science* **80**, 413–433.
- Huttegger, S. M. (2014). Learning experiences and the value of knowledge. *Philosophical Studies* **171**, 279–288.
- Jaynes, E. T. (2003). *Probability Theory. The Logic of Science*. Cambridge: Cambridge University Press.
- Jeffrey, R. C. (1957). Contributions to the theory of inductive probability. PhD dissertation, Princeton University.
- Jeffrey, R. C. (1965). *The Logic of Decision*. New York: McGraw-Hill. 3rd revised edition Chicago: University of Chicago Press, 1983.
- Jeffrey, R. C. (1968). Probable knowledge. In Lakatos, I., editor, *The Problem of Inductive Logic*, pp. 166–180. Amsterdam: North-Holland.
- Jeffrey, R. C. (1987). Alias Smith and Jones: The testimony of the senses. *Erkenntnis* **26**, 391–399.
- Jeffrey, R. C. (1988). Conditioning, kinematics, and exchangeability. In Skyrms, B. & Harper, W. L., editors, *Causation, Chance, and Credence*, Volume 1, pp. 221–255. Dordrecht: Kluwer.
- Jeffrey, R. C. (1992). *Probability and the Art of Judgement*. Cambridge: Cambridge University Press.
- Joyce, J. M. (2007). Epistemic deference: The case of chance. *Proceedings of the Aristotelian Society* **107**, 187–206.
- Joyce, J. M. (2010). The development of subjective Bayesianism. In Gabbay, D. M., Hartmann, S., & Woods, J., editors, *Handbook of the History of Logic, Vol 10: Inductive Logic*, pp. 415–476. Elsevier.
- Kadane, J. B., Schervish, M. J., & Seidenfeld, T. (1996). Reasoning to a foregone conclusion. *Journal of the American Statistical Association* **91**, 1228–1236.
- Kalai, E., & Lehrer, E. (1994). Weak and strong merging of opinions. *Journal of Mathematical Economics* **23**, 73–86.
- Kelly, T. (2005). The epistemic significance of disagreement. *Oxford Studies in Epistemology* **1**, 167–196.
- Kelly, T. (2008). Disagreement, dogmatism, and belief polarization. *Journal of Philosophy* **105**, 611–633.
- Lange, M. (2000). Is Jeffrey conditionalization defective in virtue of being non-commutative? Remarks on the sameness of sensory experiences. *Synthese* **93**, 393–403.

- Lehrer, K., & Wagner, C. G. (1981). *Rational Consensus in Science and Society: A Philosophical and Mathematical Study*. Dordrecht: D. Reidel.
- Leitgeb, H., & Pettigrew, R. (2010). An objective justification of Bayesianism II: The consequences of minimizing inaccuracy. *Philosophy of Science* **77**, 236–272.
- Levi, I. (1980). *The Enterprise of Knowledge*. Cambridge MA: MIT Press.
- Levi, I. (1987). The demons of decision. *The Monist* **70**, 193–211.
- Maher, P. (1992). Diachronic rationality. *Philosophy of Science* **59**, 120–141.
- Mongin, P. (1995). Consistent Bayesian aggregation. *Journal of Economic Theory* **66**, 313–351.
- Peirce, C. (1997). *The Fixation of Belief*, pp. 7–25. New York: Vintage Books. Originally published in 1878.
- Purves, R. A., & Sudderth, W. D. (1976). Some finitely additive probability. *The Annals of Probability* **4**, 259–276.
- Savage, L. J. (1954). *The Foundations of Statistics*. New York: Dover Publications.
- Schervish, M. J., & Seidenfeld, T. (1990). An approach to consensus and certainty with increasing information. *Journal of Statistical Planning and Inference* **25**, 401–414.
- Seidenfeld, T. (1985). Calibration, coherence, and scoring rules. *Philosophy of Science* **52**, 274–294.
- Seidenfeld, T. (1987). *Entropy and Uncertainty*, pp. 259–287. Dordrecht: D. Reidel.
- Seidenfeld, T. (2001). Remarks on the theory of conditional probability: Some issues of finite versus countable additivity. In Hendricks, V. F., editor, *Probability Theory*, pp. 167–178. Kluwer.
- Seidenfeld, T., Kadane, J. B., & Schervish, M. J. (1989). On the shared preferences of two Bayesian decision makers. *Journal of Philosophy* **86**, 225–244.
- Seidenfeld, T., Schervish, M. J., & Kadane, J. B. (2014). Non-conglomerability for countably additive measures that are not κ -additive. Manuscript CMU.
- Skyrms, B. (1985). Maximum entropy inference as a special case of conditionalization. *Synthese* **63**, 55–74.
- Skyrms, B. (1987a). Dynamic coherence. In MacNeill, B. & Umphrey, G., editors, *Advances in the Statistical Sciences VII. Foundations of Statistical Inference*, pp. 233–243. Dordrecht: D. Reidel.
- Skyrms, B. (1987b). Dynamic coherence and probability kinematics. *Philosophy of Science* **54**, 1–20.
- Skyrms, B. (1990). *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard University Press.
- Skyrms, B. (1997). The structure of radical probabilism. *Erkenntnis* **45**, 285–297.
- Talbott, W. (1991). Two principles of Bayesian epistemology. *Philosophical Studies* **62**, 135–150.
- Teller, P. (1973). Conditionalization and observation. *Synthese* **26**, 218–258.
- van Fraassen, B. C. (1980). Rational belief and probability kinematics. *Philosophy of Science* **47**, 165–187.
- van Fraassen, B. C. (1984). Belief and the will. *Journal of Philosophy* **81**, 235–256.
- van Fraassen, B. C. (1995). Belief and the problem of Ulysses and the Sirens. *Philosophical Studies* **77**, 7–37.
- van Inwagen, P. (1996). Is it wrong, everywhere, always, and for anyone, to believe anything upon insufficient evidence? In Jorand, J. & Howard-Snyder, D., editors, *Faith, Freedom, and Rationality*, pp. 137–153. Rowman & Littlefield Publishers.

Wagner, C. G. (2002). Probability kinematics and commutativity. *Philosophy of Science* **69**, 266–278.

Zabell, S. L. (2002). It all adds up: The dynamic coherence of radical probabilism. *Philosophy of Science* **69**, 98–103.

SIMON M. HUTTEGGER
DEPARTMENT OF LOGIC AND PHILOSOPHY OF SCIENCE
UNIVERSITY OF CALIFORNIA, IRVINE
SOCIAL SCIENCE PLAZA A
IRVINE, CA-92697, USA
E-mail: shuttegg@uci.edu