**The Role of Absolute Continuity in “Merging
Opinions” and “Rational Learning”**

by

Ronald I. Miller, Columbia University
Chris William Sanchirico, Columbia University

November 1996

Discussion Paper Series No. 9697-07

dp 9697-07
pages: 33

THE ROLE OF ABSOLUTE CONTINUITY IN “MERGING OF OPINIONS”
AND “RATIONAL LEARNING”

By Ronald I. Miller and Chris William Sanchirico

November 1996

Columbia University
Department of Economics
420 West 118th Street
New York, New York 10027
rm170@columbia.edu
cs282@columbia.edu

Two agents with different priors watch a sequence unfold over time, updating their priors about the future course of the sequence with each new observation. Blackwell and Dubins (1962) show that the agents' opinions about the future will converge if their priors over the sequence space are absolutely continuous: i.e., if they agree on what events are possible. From this Kalai and Lehrer (1993) conclude that the players in a repeated game will eventually agree about the future course of play and thus that “rational learning leads to Nash equilibrium.” We provide an alternative proof of convergence that clarifies the role of absolute continuity and in doing so casts doubt on the relevance of the result. From the existence of continued disagreement we construct a sequence of mutually favorable, uncorrelated “bets.” By a law of large numbers, both agents are sure that they win these bets on average over the long run and this disagreement over what is possible violates absolute continuity.

Keywords: Bayesian Learning, Learning in Games, Merging of Opinions, Rational Learning

Suppose two agents are watching a sequence of numbers unfold over time. Each begins with a prior belief over the full sequence and then updates her beliefs on the future course of the sequence as each successive coordinate is revealed. Under what conditions will the agents' posteriors on the future of the sequence converge?

In a foundational and often cited work Blackwell and Dubins (1962) show that convergence will obtain if agents' priors are absolutely continuous. More precisely, if the probability measures P and Q over infinite sequences represent the agents' priors and Q is absolutely continuous with respect to P —meaning that $QE > 0$ implies $PE > 0$ for all events E in the sequence space—then with Q probability one, the conditional distributions of P and Q on the future given the past (as defined within) will converge under the usual distance metric for measures.

More recently, Kalai and Lehrer (1993) apply this result to construct the currently preeminent model in which learning leads to Nash equilibrium. There, Blackwell and Dubins' convergence result shows that players in an infinitely repeated game will eventually “learn” to predict the future course of play through the process of continually updating their priors based on the history of play. This generates the common understanding of strategic intention necessary to obtain Nash equilibrium. In this setting the absolute continuity assumption means that each player's prior must only assign zero probability to an event when it in fact receives zero probability given players' actual strategies.

Since the absolute continuity assumption is critical for two fundamental results, a clear understanding of its precise role and interpretation is imperative. On one level, the assumption is easy to state and easier to understand: if Q thinks an event is possible, P does also. But

behind that simplicity lies a subtle idea, difficult to interpret behaviorally. Our object in this paper is to provide an alternative proof of Blackwell and Dubins' main theorem that makes transparent the role of the absolute continuity assumption. In doing so we hope to convince the reader that the result is of less relevance and efficacy than it may at first appear to be.

Blackwell and Dubins (and Kalai and Lehrer) show convergence directly in an elegant but, we feel, unintuitive proof using Radon-Nikodym and Martingale Convergence. Our approach is less elegant, but also less mysterious. We prove Blackwell and Dubins by contradiction, employing no machinery more sophisticated than a law of large numbers. The basic idea of our proof is as follows. Non-convergence means that Q thinks there is some chance that P and Q will infinitely often disagree about the probability of future events by at least some fixed ε . Each time Q and P so disagree they can make a (zero-sum) "bet" that both expect to win. An infinite number of disagreements means an infinite sequence of bets. The law of large numbers then implies that (conditional on infinite ε -disagreement), Q is certain that she will win on average in the long run. P on the other hand, will be just as sure of the opposite. We have thus translated a persistent ε -disagreement into a stark disagreement about what is and is not possible, and this violates absolute continuity.

Suppose, for example, that our agents observe a binary sequence. Both agents are sure that the sequence is generated by i.i.d. Bernoulli random variables, but agent P thinks the Bernoulli probability is $\frac{1}{2}$, while agent Q thinks the probability is $\frac{1}{2} + \varepsilon$. Clearly, P and Q do not converge in the sense of Blackwell and Dubins. Where then is the event defeating Q 's absolute continuity with respect to P ? In this simple example, Q is certain that the average

frequency of ones in the Bernoulli trials will converge to $\frac{1}{2} + \varepsilon$, while P , who thinks it must converge to $\frac{1}{2}$, regards this as impossible.

This same idea can be restated in terms of the betting discussed above. Q always thinks it ε -more likely than P that the *next* element is a one. We can then structure a bet in which both expect to win a dollar in every period: Q pays P x if the next element is a one and y (negative) if the next element is a zero where x and y solve the linear equations for the agents' expectations:

$$\begin{aligned}\frac{1}{2}x + \frac{1}{2}y &= 1 \\ (\frac{1}{2} + \varepsilon)x + (\frac{1}{2} - \varepsilon)y &= -1\end{aligned}\tag{1}$$

Applying the law of large numbers, Q believes *for sure* that P will pay her one on average over the long term, while P thinks there is no chance of this.

Another way to understand this example, and the strength of absolute continuity generally, is in terms of convergent statistics. Absolute continuity says that whenever Q thinks an event is possible, P must agree. Thus any sequence of random variables that converges to some value P -almost everywhere must converge to the same value Q -almost everywhere. In a sense then, P and Q must agree on all P -a.s. convergent “statistics.” This is an extremely strict requirement because the entire universe of all conceivable P -convergent statistics is on trial. In this example, the average number of Bernoulli ones, the mean of the sequence, is a P -convergent statistic upon which P and Q disagree.

In this paper we show how the betting approach generalizes to arbitrary probabilities. In effect we show that it is always possible to tailor a convergent statistic to exploit any persistent

disagreement between P and Q regarding the continuation of the game. Importantly, the statistic is quite intuitively related to the disagreement to which it was tailored. It is essentially P 's time-average winnings in a non-overlapping sequence of bets made against Q when the two disagree by at least ε . Since both players always expect to win each bet, they must both be sure that they will win on average over the long run. The main text constructs bets for the simplest case in which each factor of the sequence space is finite and both P and Q place positive measure on every partial history. An appendix shows that our construction generalizes to countably generated coordinate spaces, such as $\mathcal{R}^n$ with the Borel sets.

Before moving forward with our main argument, it is worth pausing to switch to an alternate perspective that helps illustrate the inherent implausibility of absolute continuity and convergence itself. Both the assumption of absolute continuity and the convergence theorem concentrate on the “global” measure that assigns probabilities to an uncountable number of infinite-length histories. It is easier to understand convergence from a “local” perspective. This entails first viewing histories as paths through an infinite length tree and partial histories as nodes in this tree and next, viewing probabilities on this tree as collections of “local” measures, one for each node. These local measures tell us the probability of each branch emanating from the node. We constructed the priors in our example locally. Since the measures there were i.i.d., all nodes had the same local measure: e.g., $(\frac{1}{2}, \frac{1}{2})$ for P . The local and global perspectives are equivalent (for the measures considered by Blackwell and Dubins¹). Bayes' rule is merely the “isomorphism” that translates the global language of priors into the local language of collections of nodal measures. Once Bayes' rule is seen in this light, its apparent power for generating learning is much diminished.

Indeed, from the local perspective, there is little reason to expect merging of opinions, which would require our two agents come to agree on the likelihood of the current draw. Speaking informally, suppose we constructed P by randomly assigning a local measure to each node and then repeated the experiment to construct Q . Why would we ever expect the collections of local measures to begin to look the same for nodes far enough out in the tree? Absolute continuity is thus much more than a mere regularity condition; it imposes a great deal of structure on the priors. Our object in this paper is show how such a seemingly plausible assumption produces such implausible results.

1. FRAMEWORK

Let each $X(i)$, $i = 1, 2, \dots$ be a finite set.² Let $X = \prod_{i=1}^{\infty} X(i)$ be the set of *(full) histories*, $X_n = \prod_{i=1}^n X(i)$ the set of *(partial) histories (up to n)* and $X'' = \prod_{i=n+1}^{\infty} X(i)$, the set of *continuations (from n)*. It is helpful to think of X as the set of paths through an infinite tree. X_n would then represent the set of all nodes of rank n , and X'' the subtree following each node in X_n . We endow each of X, X_n, X'' with their usual product σ -algebras. In particular, write $\mathcal{B}$, $\mathcal{B}_n$ and $\mathcal{B}''$ for the σ -algebras on X , X_n and X'' , respectively. ($\mathcal{B}_n$ is trivial but it is helpful to have the notation.) We denote the typical partial history as $h_n \in X_n$. The partial history of a given full history x is written $h_n(x)$, which is just the projection map $h_n: X \rightarrow X_n$.

For any measures R and S on a given measurable space $(Y, \mathcal{A})$, define the *distance* $|R - S| = \sup_{A \in \mathcal{A}} |R(A) - S(A)|$. We say that R is *absolutely continuous* with respect to S , written $R \ll S$, if $R(A) > 0$ implies $S(A) > 0$. Let P and Q be probability measures on X such that³ $P(h_n), Q(h_n) > 0$ for all⁴ h_n . Thus we may define for each history, h_n , P 's *conditional distribution on the future* X^n given the past h_n , $P^n(h_n): \mathfrak{E}^n \rightarrow \mathfrak{R}$, s.t.

$$P^n(h_n)(E) = P(X_n \times E | h_n), \text{ for all events } E \text{ on } X^n.$$

Define $Q^n(h_n)$ similarly. In terms of a tree, $P^n(h_n)$ and $Q^n(h_n)$ are measures on the subtree following node h_n induced by P and Q .

For the convenience of the reader, we restate Blackwell and Dubins' convergence theorem for the special case we are considering.⁵

Theorem 1 [Blackwell and Dubins' (1962)]: If Q is absolutely continuous with respect to P , then for every history $x \in X$ in some set of Q probability 1, the distance

$|P^n(h_n(x)) - Q^n(h_n(x))|$ between P and Q 's conditional distributions on the future given the past converges to zero as n goes to infinity.

2. CONTRADICTING CONVERGENCE

The first step in our proof by contradiction is to illuminate what it means for the conditional probabilities of P and Q to *not* converge in the sense of Blackwell and Dubins. Simply contradicting Theorem 1 yields: *for each infinite history x in some set of positive Q measure there exists an $\varepsilon_x > 0$ such that for infinitely many n , we can find a continuation*

event $D^n \in \mathcal{E}^n$ with $|P^n(h_n(x))(D^n) - Q^n(h_n(x))(D^n)| > \varepsilon_x$. In other words, for each such

x there is *persistent ε_x -disagreement*. Let $\mathcal{C}^n$ represent the collection of (finite) cylinders in

the continuation X^n : sets of the form $A \times X^m \subseteq X^n$, where $m > n$ and $A \subseteq \prod_{i=n+1}^m X(i)$.

Standard arguments establish that we lose no generality if, in contradicting convergence, we 1) take ε_x uniformly over x and 2) take the continuation events on which P and Q disagree to be cylinders. This is stated formally in the following lemma, the proof of which is relegated to an appendix.

Lemma 1 [Contradicting Convergence]: Suppose that along some set of histories with positive Q measure $|P^n(h_n(x)) - Q^n(h_n(x))|$ does not converge to zero. Then there exists an $\varepsilon > 0$ and a set D of positive Q measure, such that along all infinite histories x in D there are infinitely many times n at which we can find a cylinder $C^n \in \mathcal{C}^n$ in the continuation from n with $|P^n(h_n(x))(C^n) - Q^n(h_n(x))(C^n)| > \varepsilon$.

The two components of this lemma correspond to the two conditions that we will require to apply a law of large numbers. The fact that cylinder events resolve in finite time will allow us to construct a sequences of bets which do not overlap in time and are thus uncorrelated. The uniformity in ε allows a uniform bound on the “stakes” of the bets.

3. DEFINING THE SEQUENCE OF BETS

From the infinite sequence of ε -disagreements on cylinders along each path in D , we construct a sequence of uncorrelated zero-sum “bets” each of which P and Q both expect to

win. A law of large numbers translates this persistent difference in *expectation* into a difference of opinion about what long run average outcomes are *possible*. Differing views about what is possible rule out absolute continuity.

The construction of these bets requires a fair amount of notation, but is quite intuitive, and can be seen as a three-part generalization of the simple example in the introduction. The first generalization is to bets that are uncorrelated, as opposed to independent. This poses no problem as we can simply apply a different law of large numbers. The second generalization accounts for the fact that infinite ε -disagreement occurs only along a *subset* of paths—this turns out to be just a matter of bookkeeping. The last generalization allows for the possibility that bets take different amounts of time to resolve. To deal with this we convert from calendar time, n , to an event time, k , based on the resolution times of the bets.

We will say that *there is a (ε) -bet at node h_n* when $|P^n(h_n) - Q^n(h_n)| > \varepsilon$. For all nodes h_n that have bets, we arbitrarily select (by the axiom of choice) a cylinder event $C_{h_n} \in \mathcal{C}^n$ satisfying $|P^n(h_n)(C_{h_n}) - Q^n(h_n)(C_{h_n})| > \varepsilon$ and define the *bet at node h_n* to be the following random variable on the continuation: $B_{h_n}: X^n \rightarrow \mathbb{R}$, s.t.

$$B_{h_n}(y) = \frac{1}{p-q} \times \begin{cases} 2-(p+q), & y \in C_{h_n} \\ -(p+q), & y \notin C_{h_n} \end{cases},$$

where we abbreviate as $p = P^n(h_n)(C_{h_n})$ and $q = Q^n(h_n)(C_{h_n})$. By convention, if there is no bet at h_n , we set $B_{h_n} \equiv 0$. We chose these two particular values for B_{h_n} because they imply that the expected value of B_{h_n} is one under $P^n(h_n)$ and is negative one under $Q^n(h_n)$.

Where there is ε -disagreement we think of B_{h_n} as a bet between P and Q as a result of their (conditional) disagreement about whether C_{h_n} will occur in the continuation. The realized value of B_{h_n} is the net payment from Q to P . Suppose, for example, that P thinks C_{h_n} is at least ε more likely than does Q . The bet says that Q pays P the amount $\frac{2-(p+q)}{p-q}$ if C_{h_n} occurs and “pays” the negative amount $-\frac{p+q}{p-q}$ if not, and each player expects to win one dollar. Importantly, since $|p-q| > \varepsilon$ the stakes of all bets are uniformly bounded across all nodes $h_n: |B_{h_n}| \leq \frac{2}{\varepsilon}$.

We have defined bets at each node where it is possible: next we assemble a sequence of bets so that the bet upon cylinders do not overlap in time. This will guarantee that the sequence of bets is uncorrelated (in a particular sense) so that we can apply a law of large numbers. Define $n_1(x)$, the *time of the first bet* to be the date of the first node along x that has a bet. If there is no bet at any node along x , we set $n_1(x) = \infty$. We now “patch together” the nodal bets described above to obtain the first bet in our constructed sequence. Let $c_n(x)$ be the continuation of history x , i.e. the projection of x onto X^n . If $n_1(x) = \bar{n} < \infty$, set the *first bet* $B_1(x) = B_{h_{\bar{n}}}(c_{\bar{n}}(x))$. If $n_1(x) = \infty$ we let $B_1(x) = 0$. Thus, for any history x along which there is *some* bet, $B_1(x)$ is the outcome of the *first* such bet: i.e. the net payment from Q to P . If there is no bet along x , $B_1(x) = 0$.

Recall that we have associated a particular cylinder in the continuation with each node having a bet. As a cylinder this event may be expressed as $A_n^r \times X^{r+1}$, for some r , where

$A'_n \subseteq X(n+1) \times \dots \times X(r)$. If h_n has a bet, the *resolution time* r_{h_n} of the bet at h_n is defined to be the smallest such r . The *resolution time*, $r_1(x)$, of the first bet is r_{h_1} at the first betting node, $h_1(x)$ along x . If there is no first betting node along x (i.e., $n_1(x) = \infty$), set $r_1(x) = \infty$.

In order to complete an inductive definition, suppose that $r_{k-1}(x), n_{k-1}(x), B_{k-1}(x)$ have already been defined. The *time of the k^{th} bet*, $n_k(x)$, will be the smallest n such that $n \geq r_{k-1}(x)$ and there is a bet at $h_n(x)$. Set $n_k(x) = \infty$ if there is no such n , either because $r_{k-1}(x) = \infty$ or because there is no bet at any node along x following time $r_{k-1}(x)$. Then define $B_k(x)$ and $r_k(x)$ in the same way as $B_1(x)$ and $r_1(x)$.

It should be noted that along all histories x in the set D of persistent ε -disagreements, as described in Lemma 1, there is a non-degenerate bet for all k . Since there are infinite ε disagreements along every history $x \in D$, we can always find yet another point of disagreement after the resolution of the last bet. We do not use *all* of the infinity of bets along history x ; just a subsequence constructed so that the $k-1^{\text{th}}$ bet is resolved before the k^{th} bet is made.

4. APPLYING THE LAW OF LARGE NUMBERS.

We have created a sequence $B_1, B_2, \dots$ of random variables on X where each B_k describes the k^{th} bet made between P and Q . The B_k 's are not independent, because the bet upon event and its probabilities under P and Q depend on the history to date. However, they are essentially uncorrelated: their expectation (conditional on being non-degenerate) is the same across all histories up to k (one for P and negative one for Q); in particular, it is the same

across all realizations of $B_1, \dots, B_{k-1}$. A slight modification of the B_k 's will make them truly uncorrelated.

For all k , let B_k^P and B_k^Q be defined from B_k by setting $B_k^P(x) = B_k^Q(x) = B_k(x)$, wherever $B_k(x) \neq 0$ and $B_k^P(x) = -B_k^Q(x) = 1$, wherever, $B_k(x) = 0$. Note that on the set D of infinite ε -disagreements, $B_k^P = B_k^Q = B_k$ for all k .

As is well known, random variables Z and Y are uncorrelated if the expected value of Z conditional on Y is constant across values for Y : i.e. $E[Z|Y=y]$ is constant across y ; for then $E[ZY] = \sum_y \Pr(Y=y)E[ZY|Y=y] = \sum_y y \Pr(Y=y)E[Z|Y=y] = E[Z]E[Y]$. Thus since the k^{th} bet is made after the $k-j^{\text{th}}$ bet is resolved and the expectation of the k^{th} bet from when it is made is always the same, the B_k^P are P -uncorrelated: $E_P[B_k^P | B_{k-j}^P = b] = 1$, for all b . Similarly, the B_k^Q are Q -uncorrelated.⁶ Further, since for each h_n , $|B_{h_n}| \leq \frac{2}{\varepsilon}$ the same uniform bound applies to both B_k^P and B_k^Q , for all k . These two conditions, orthogonality and uniform boundedness, are more than enough to apply a law of large numbers⁷ (see, e.g. Stokey and Lucas⁸ (1989), p. 422) This enables us to conclude that $\frac{1}{K} \sum_{k=1}^K B_k^P$ converges to one P -almost surely and $\frac{1}{K} \sum_{k=1}^K B_k^Q$ converges to negative one Q almost surely. Now consider the event that we are in D and P 's average net winnings $\frac{1}{K} \sum_{k=1}^K B_k$ go to -1: i.e. the event $D \cap \left\{ \frac{1}{K} \sum_{k=1}^K B_k \rightarrow -1 \right\}$. Since $B_k^P = B_k^Q = B_k$ on D , and $Q(D) > 0$, Q thinks this event is possible (indeed certain, conditional on D):

$$Q\left(D \cap \left\{\frac{1}{K} \sum_{k=1}^K B_k \rightarrow -1\right\}\right) = Q(D) > 0.$$

P on the other hand regards this event as impossible since she is certain that her average winnings converge to one on D :

$$P\left(D \cap \left\{\frac{1}{K} \sum_{k=1}^K B_k \rightarrow 1\right\}\right) = P(D) \Rightarrow P\left(D \cap \left\{\frac{1}{K} \sum_{k=1}^K B_k \rightarrow -1\right\}\right) = 0.$$

We conclude that Q is not absolutely continuous with respect to P .

5. NOTES AND EXAMPLES

5.1 Countable Intuition in an Uncountable Setting

Perhaps the reason that the absolute continuity assumption has gained such currency in the literature is that it is so plausible in a finite, or even countable setting. Even the stronger assumption that both players regard each state as at least *possible* seems attractive, since all this rules out is dogmatism. But it would be a mistake to carry this intuition into the necessarily uncountable setting that is relevant here: obviously, in this case some events must receive zero measure.

In the finite case *every* measure is absolutely continuous with respect to the uniform measure, while this is far from the case in the uncountable setting. The measure built from nodal probabilities of $(\frac{1}{2}, \frac{1}{2})$ in our introductory example, for instance, is the uniform measure on binary sequences. And indeed we used it as an example because it is trivial to find a measure that is not absolutely continuous with respect to it. Thus, the uniform measure, which

we think of as so “open-minded” in the finite case is highly “opinionated” in the uncountable setting.

5.2 Countable Restrictions in X

It has been proposed in several contexts that Blackwell and Dubins’ convergence result be applied to the special case of measures that have the same countable support in X . Here absolute continuity is satisfied if both P and Q put positive weight on all histories in this common support, which is only possible since the support is countable. *Given this set-up* absolute continuity seems to be a reasonable assumption; but the problem has just been transferred to the set-up itself. For this approach begs the question of which set should serve as common support. How should we suppose that the players in a repeated game, for instance, are able to agree a priori on one set of strategies from among the uncountable multitude of such subsets? The very assumption of so much initial agreement accomplishes most of the learning by itself.

In the case of repeated games, the countable restriction is accomplished if players understand that all strategies are drawn from some given countable set. In this case, the resort to countable supports is particularly unpromising. For as Nachbar (1996) shows there is no guarantee that the best responses to any given countable subset will lie in that set.

5.3 The Grain of Truth.

Kalai and Lehrer support the assumption of absolute continuity with reference to a stronger assumption called the “the grain of truth.” In our context this means that P may be written as a convex combination of Q and any other measure on X . (Note that the functional convex combination of two probability measures is as well a probability measure.)

Interpreting Q as the “true” measure on X and P as an observer’s beliefs, the observer’s beliefs must contain a “grain of truth.”

Our objection to this is very closely related to our objection to the countable restriction discussed above. For if there were only a countable number of probability measures under consideration, not assigning at least some small probability to each would seem dogmatic. But when, as here, the number of possible measures on X is uncountable, we run into the same problem as above. In order for the grain of truth criterion to hold, we have to know that the true measure lies within a particular countable subset of the uncountable number of possibilities. This means that we must know a great deal even before observing the first element of the true sequence.⁹

6. CONCLUSION

We have provided a proof of Blackwell and Dubins’ convergence result that demystifies the role of absolute continuity in a way that we believe makes the result less compelling. Essentially we have provided an interpretation of the lack of absolute continuity *in terms of* persistent disagreement. Lacking any independent behavioral interpretation, absolute continuity becomes the economic and behavioral equivalent of convergence.

In general, we feel that there is little promise in attempting to derive convergence from “regularity” conditions on priors. A shift to the “local perspective” laid out in the introduction makes clear that convergence simply does not follow from the structure of the problem. We do not conclude from this, however, that the research program started by Blackwell and Dubins and then continued by Kalai and Lehrer should be abandoned, only that

its focus should shift. The most fruitful approach, in our view, would be to recognize that results require genuine, substantive restrictions on beliefs and then propose restrictions that are explicitly grounded in our best understanding of human behavior.

7. APPENDIX: PROOF OF LEMMA 1

The fact that ε may be chosen uniformly follows from standard arguments and the continuity of probability measures. Let D' be the positive Q measure set of infinite histories, x , with an ε_x and infinitely many n at which we can find a continuation event $D'' \in \mathcal{B}^n$ with

$|P^n(h_n(x))(D'') - Q^n(h_n(x))(D'')| > \varepsilon_x$. Let D_j be the subset of D' such that for infinitely many n , we can find a continuation event $D'' \in \mathcal{B}^n$ with $|P^n(h_n(x))(D'') - Q^n(h_n(x))(D'')| > \frac{1}{j}$.

Then $D_j \uparrow D'$ and so by continuity of probabilities, $Q(D') > 0$ implies $Q(D_j) > 0$ for large enough j . Therefore, for some j large enough, we can set $D = D_j$ and $\varepsilon = \frac{1}{j}$.

It remains to show that we may restrict attention to disagreements over cylinders. This follows from the fact that the distance norm used by Blackwell and Dubins is equivalently defined by taking the supremum only over a generating algebra (here the cylinders): that is, if the algebra $\mathcal{F}_0$ generates the σ -algebra $\mathcal{F}$, in some measurable space $(\Omega, \mathcal{F})$, then

$$|P - Q| \equiv \sup_{S \in \mathcal{F}} |P(S) - Q(S)| = \sup_{A \in \mathcal{F}_0} |P(S) - Q(S)|. \text{ This, in turn, may be deduced directly from}$$

Halmos' Monotone Class Theorem (Billingsley (1995), p. 43), or alternatively from a well-known "Approximation Theorem" (See, e.g., Ash p. 20): for every $\varepsilon > 0$ and every $S \in \mathcal{F}$, there exists $A \in \mathcal{F}_0$ s.t. $P(S \Delta A) < \varepsilon$, where $S \Delta A = (S - A) \cup (A - S)$. ■

8. APPENDIX: GENERALIZATION TO UNCOUNTABLE COORDINATE SPACES (WITH COUNTABLY GENERATED σ -ALGEBRAS)

The logic for the uncountable case is the same as for the finite case analyzed in the body of the paper. The only difficulty is insuring that the constructed bets B_k are measurable functions. For the finite case we simply selected a bet at each partial history having one, and from these “nodal bets” we constructed the B_k sequence. In the uncountable case, we can still make such a selection, by the axiom of choice. The problem is that the selection must be made in such manner that the bets B_k are $\mathcal{B}$ -measurable functions. Here we show how this is possible when each coordinate σ -algebra is countably generated—as the Borel sets on $\mathfrak{R}$, $\mathfrak{R}^n$, or even $\mathfrak{R}^\infty$.

As in Blackwell and Dubins, we restrict attention to probabilities on $(X, \mathcal{B})$ that can be constructed from measures on the future conditional on each partial history. To state this restriction formally, we restate their

Definition 1 [Blackwell and Dubins (1962)]: A probability measure P on $(X, \mathcal{B})$ is *predictive* if for every $n \geq 1$, there exists a *conditional distribution* P^n for the future X^n given the past; that is, if there exists a function $P^n(x_1, \dots, x_n)(E)$ where $x_1, \dots, x_n$ ranges over X_n and E ranges over $\mathcal{B}^n$ with the usual three properties: $P^n(x_1, \dots, x_n)(E)$ is $\mathcal{B}_n$ -measurable for fixed E ; is a probability distribution on $(X^n, \mathcal{B}^n)$ for fixed $(x_1, \dots, x_n)$; and for bounded $\mathcal{B}$ -measurable ϕ

$$\int \phi \, dP = \int \left(\int \phi(x_1, \dots, x_n, z_{n+1}, z_{n+2}, \dots) \, dP^n(x_1, \dots, x_n) \right) dP_n, \quad (2)$$

where P_n is the *marginal distribution* of P on $(X_n, \mathcal{B}_n)$; that is, $P_n(A) = P(A \times X^n)$, for all $A \in \mathcal{B}_n$.

Though not all probability spaces are predictive, any probability measure on the Borel Sets with respect to a complete separable measure space will be (See e.g. Ash p. 266, paragraph 2).

Next we restate Blackwell and Dubins' convergence result:

Theorem 2 [Blackwell and Dubins' "Main Theorem"]: Suppose that P is predictive on $(X, \mathcal{B})$ and that Q is absolutely continuous with respect to P . Then for each conditional distribution P^n of the future given the past with respect to P , there exists a conditional distribution Q^n of the future given the past with respect to Q such that, with the exception of a set of histories $(x_1, \dots, x_n, x_{n+1}, \dots)$ of Q -probability 0, the distance between $P^n(x_1, \dots, x_n)$ and $Q^n(x_1, \dots, x_n)$ converges to 0 as n converges to ∞ .

We add an additional assumption to prove a stronger result:

Assumption 1: Each coordinate $\mathcal{B}(i)$ $i = 1, 2, \dots$ is countably generated.^{10,11}

Theorem 3: Suppose that P is predictive on $(X, \mathcal{B})$ and that Q is absolutely continuous with respect to P . Then for each conditional distribution P^n of the future given the past with respect to P and all conditional distributions Q^n of the future given the past with respect

to Q , with the exception of a set of histories $(x_1, \dots, x_n, x_{n+1}, \dots)$ of Q -probability 0, the distance between $P^n(x_1, \dots, x_n)$ and $Q^n(x_1, \dots, x_n)$ converges to 0 as n converges to ∞ .

The assumption that $\mathcal{B}(i)$ is countably generated is weaker than it may at first seem: as noted above the Borel sets on $\mathcal{R}^n$, even $\mathcal{R}^\infty$ are countably generated. Thus, generation by a countable collection of sets is a much weaker condition than generation by a countable partition, the assumption imposed by Kalai and Lehrer (1993) in their proof of Blackwell and Dubins type convergence.

8.1 Contradicting Blackwell and Dubins Convergence within a Countable Collection of Continuation Events

Let $\mathcal{F}(i)$ be the smallest algebra containing the countable collection of subsets generating $\mathcal{B}(i)$. $\mathcal{F}(i)$ will also be countable. Let $\mathcal{F}^n$ be the collection of all finite disjoint unions of sets that may be expressed as $F(n+1) \times \dots \times F(m) \times X(m+1) \times \dots$ for any $m > n$ and any $F(j) \in \mathcal{F}(i)$ all $j = n+1, \dots, m$. By standard techniques:

Lemma 2: $\mathcal{F}^n$ is a countable algebra on the continuation space X^n and generates the continuation σ -algebra $\mathcal{B}^n$.

By the same reasoning as in *Lemma 1* we obtain the following restatement of non-convergence:

Lemma 3: Suppose that for all x on some set of Q positive measure, there exists $\varepsilon_x > 0$ and infinitely many n such that we can find a continuation event $D^n \in \mathcal{E}^n$ satisfying

$$\left| P^n(x_1, \dots, x_n)(D^n) - Q^n(x_1, \dots, x_n)(D^n) \right| > \varepsilon_x. \text{ Then there exists an } \varepsilon > 0 \text{ such that for all } x \text{ on}$$

some set D of Q positive measure there exist infinitely many n such that we can find an

$$\text{event } F \in \mathcal{F}^n \text{ such that } \left| P^n(x_1, \dots, x_n)(F) - Q^n(x_1, \dots, x_n)(F) \right| > \varepsilon.$$

8.2 Structuring the Bets

8.2.1 A Bet for Each Event in $\mathcal{F}^n$ at each Partial History $x_1, \dots, x_n$

The first step is to define a bet for each time n , partial history $x_1, \dots, x_n$ and each element of $\mathcal{F}^n$ on which P and Q disagree by at least ε . For each $n = 1, 2, \dots$, enumerate $\mathcal{F}^n$ as $\{F_m^n\}$. Set $P_m^n: X \rightarrow \mathfrak{R}$ s.t. $P_m^n(x) = P^n(h_n(x))(F_m^n)$. Define Q_m^n in the same manner. The functions P_m^n and Q_m^n are $\mathcal{E}_n$ -measurable by Definition 1 and the measurability of the projection h_n .

For all $n = 1, 2, \dots$, $m = 1, 2, \dots$ define the *random bet on the m^{th} event at time n* ,

$$B_m^n: X \rightarrow \mathfrak{R} \text{ s.t.}$$

$$B_m^n = \begin{cases} I_{X_n \times F_m^n} \frac{2 - (P_m^n + Q_m^n)}{P_m^n - Q_m^n} + I_{X_n \times \neg F_m^n} \frac{-(P_m^n + Q_m^n)}{P_m^n - Q_m^n}, & \text{if } |P_m^n - Q_m^n| > \varepsilon \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

To aid interpretation we will say that *there is a bet at partial history* $x_1, \dots, x_n$ (alternatively *along history* x *at time* n) *on the event* F_m^n if $|P_m^n(x) - Q_m^n(x)| > \varepsilon$. (Note the slight abuse of notation.)

We will also define the (non random) *resolution time* r_m^n of the m^{th} bet at time n . Since F_m^n is the disjoint union of finite rectangles, there exists r such that F_m^n may be expressed as a union of sets of the form $F(n+1) \times \dots \times F(r) \times X(r+1) \times \dots$, where $F(i) \in \mathcal{F}(i)$, all $n < i \leq r$. Let $r_m^n > n$ be the smallest such r for F_m^n . Then $\mathcal{B}_{\hat{n}}$ where $\hat{n} = r_m^n$, which may be written $\mathcal{B}_{r_m^n}$, is the σ -algebra describing what is known about x when the m^{th} bet at time n is resolved.

Lemma 4 [Properties of the B_m^n]:

- i) All B_m^n are $\mathcal{B}_{r_m^n}$ -measurable (a fortiori $\mathcal{B}$ -measurable) functions.
- ii) If there is a bet at partial history $x_1, \dots, x_n$ on the event F_m^n , then

$$\int B_m^n(x_1, \dots, x_n, z_{n+1}, z_{n+2}, \dots) dP^n(x_1, \dots, x_n) = 1$$

$$\int B_m^n(x_1, \dots, x_n, z_{n+1}, z_{n+2}, \dots) dQ^n(x_1, \dots, x_n) = -1.$$

- iii) The B_m^n are uniformly bounded: in particular, $\forall m, n$, $|B_m^n| \leq \frac{2}{\varepsilon}$.

Proof: Clear from the definitions. ■

8.2.2 A Measurable Selection of Bets for Each Time n

The second step is to select measurably a bet for each partial history. We will arbitrarily choose the first event in our enumeration of $\mathcal{F}^n$ for which a bet exists. Thus the *random bet-upon event* for n is $m_n: X \rightarrow N \cup \{\infty\}$ s.t.

$$m_n(x) = \sup \left\{ m \in N \mid \left| P_m^n(x) - Q_m^n(x) \right| \leq \varepsilon, \forall m' < m \right\}.$$

Note that $m_n(x) = \infty$ if there's no bet at $x_1, \dots, x_n$.

Lemma 5: For all $n = 1, 2, \dots$ the random bet-upon event at time n , $m_n: X \rightarrow N \cup \{\infty\}$ is a $\mathcal{B}_n$ -measurable (a fortiori $\mathcal{B}$ -measurable) function.¹²

Proof: Fix m and n . $A_m^n \equiv \left\{ (x_1, \dots, x_n) \in X_n \mid \left| P_m^n(x_1, \dots, x_n)(F_m^n) - Q_m^n(x_1, \dots, x_n)(F_m^n) \right| > \varepsilon \right\}$ is a

$\mathcal{B}(1) \times \dots \times \mathcal{B}(n)$ -measurable subset of X_n by Definition 1. Then

$E_m^n \equiv \left\{ \left| P_m^n - Q_m^n \right| > \varepsilon \right\} = A_m^n \times X(n+1) \times \dots$ is a $\mathcal{B}_n$ -measurable set. Therefore, given n , $\forall m$,

$\{m_n = m\} = E_m^n - \bigcup_{m' < m} E_{m'}^n$ and $\{m_n = \infty\} = -\bigcup_{m'=1}^{\infty} E_{m'}^n$ are $\mathcal{B}_n$ -measurable sets. ■

8.2.3 Defining the Sequence of Bets

Lastly, we define the sequence of bets themselves. Define the *random time of the first bet* $n_1: X \rightarrow N \cup \{\infty\}$ s.t.

$$n_1(x) = \sup \left\{ n \in N \mid \forall n' < n, \forall m, \left| P_m^{n'}(x) - Q_m^{n'}(x) \right| \leq \varepsilon \right\} = \sup \left\{ n \in N \mid \forall n' < n, m_{n'}(x) = \infty \right\}.$$

Note that if there is no bet along x at time n' , then $m_{n'}(x) = \infty$.

We may now define *the first bet*, $B_1: X \rightarrow \mathbb{R}$ s.t.

$$B_1 = B_{m_{n_1}}^{n_1} = \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} I_{n_1=n \wedge m_n=m} B_m^n,$$

where $I_{n_1=n \wedge m_n=m}$ is the indicator function for the set $\{m_n(x) = m \text{ and } n_1(x) = n\}$. For each history x we have assigned at most one first bet time $n_1(x)$: all histories that have bets at some point n are assigned a betting time. Further, for each history x and its finite betting time $n_1(x) = n < \infty$, should it have one, we have assigned one bet upon event $m_n(x) < \infty$. Thus $I_{n_1=n \wedge m_n=m}$ equals one for at most one n, m pair. Thus for each history x , our summation formulation either selects a unique bet, or if there is none, equals zero.

The *random resolution time of the first bet* is defined in the same manner as B_1 ; i.e.,

$$r_1: X \rightarrow N \cup \{\infty\} \text{ s.t.}$$

$$r_1 = r_{m_{n_1}}^{n_1} = \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} I_{n_1=n \wedge m_n=m} r_m^n.$$

The *random time of the k^{th} bet* is the first time after resolution of the $k-1^{\text{th}}$ bet that a bet exists. First, set

$$n_k(x) = \sup\{n \mid n \geq r_{k-1}(x) \text{ and } \forall n' \text{ s.t. } n' \geq r_{k-1}(x) \text{ and } n' < n, m_{n'}(x) = \infty\}.$$

Then just as for $k=1$, set

$$B_k = B_{m_{n_k}}^{n_k} = \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} I_{n_k=n \wedge m_n=m} B_m^n$$

$$r_k = r_{m_{n_k}}^{n_k} = \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} I_{n_k=n \wedge m_n=m} r_m^n.$$

The next two results concern the measurability of the objects defined in this subsection.

Lemma 6 [Measurability of Bets and Times]: For all bets $k = 1, 2, \dots$ and all times

$n = 1, 2, \dots$ both $\{n_k = n\}$ and $\{r_k = n\}$ are $\mathcal{B}_n$ -measurable sets.

Proof: For all times n , define $r^n: X \rightarrow N \cup \{\infty\}$ s.t. $r^n = \sum_{m=1}^{\infty} I_{m_n=m} r_m^n$. Since m_n is $\mathcal{B}_n$ -measurable (Lemma 5), so is r^n .

$k=1$: Since for all n , $\{m_n < \infty\}$ is a $\mathcal{B}_n$ -measurable set,

$\{n_1 = n\} = \{m_n < \infty\} - \bigcup_{n' < n} \{m_{n'} < \infty\}$ is also a $\mathcal{B}_n$ -measurable set. Thus

$\{r_1 = n\} = \bigcup_{n' < n} [\{n_1 = n'\} \cap \{r^{n'} = n\}]$ is a $\mathcal{B}_n$ -measurable set.

$k-1 \rightarrow k$: Suppose that for all n , $\{r_{k-1} = n\}$ is a $\mathcal{B}_n$ -measurable set. Combining with

Lemma 5 yields that $\{n_k = n\} = \bigcup_{n' \leq n} [\{r_{k-1} = n'\} \cap [\{m_n < \infty\} - \bigcup_{n'' \leq n' < n} \{m_{n''} < \infty\}]]$ is a $\mathcal{B}_n$ -

measurable set. Therefore, $\{r_k = n\} = \bigcup_{n' < n} [\{n_k = n'\} \cap \{r^{n'} = n\}]$ is also a $\mathcal{B}_n$ -measurable

set. ■

Corollary 1: For all bets $k = 1, 2, \dots$, the betting time n_k , the resolution time r_k and the bet B_k itself are $\mathcal{B}$ -measurable functions.

As in the text, define for all $k = 1, 2, \dots$, $B_k^P: X \rightarrow \mathfrak{R}$ and $B_k^Q: X \rightarrow \mathfrak{R}$ s.t.

$$B_k^P = \begin{cases} B_k & n_k < \infty \\ 1 & n_k = \infty \end{cases}, \quad B_k^Q = \begin{cases} B_k & n_k < \infty \\ -1 & n_k = \infty \end{cases}.$$

Lemma 7 [Properties of B_k for Application of Law of Large Numbers]:

- i) B_k and so B_k^P and B_k^Q are uniformly bounded across k
- ii) $E_P[B_k^P] = 1$ and $E_Q[B_k^Q] = -1$, all $k = 1, 2, \dots$
- iii) $\text{cov}_P[B_k^P, B_j^P] = \text{cov}_Q[B_k^Q, B_j^Q] = 0$, all $j, k = 1, 2, \dots, j \neq k$,
- iv) $\{B_k\}$, $\{B_k^P\}$ and $\{B_k^Q\}$ coincide on D , the set of histories along which there are infinitely many ε -disagreements.

Proof: We consider only B_k^P ; the arguments for B_k^Q are the same. i) follows immediately from the uniform boundedness of the B_m^n . iv) is clear from the definition of B_k^P .

ii): Roughly, wherever there is an n^{th} bet, its conditional expected value is one by construction. By definition of B_k^P , the conditional expected value is also one on the set of histories which have no n^{th} bet. The following argument shows how to aggregate across the countable number of n, m pairs to obtain the unconditional expected value.

Since $\int B_k^P dP = P[n_k = \infty] + \int B_k dP$, it suffices to show $\int B_k dP = P[n_k < \infty]$. By the

bounded convergence theorem:

$$\begin{aligned} \int B_k dP &= \int \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} I_{n_k = n \wedge m_n = m} B_m^n dP = \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} \int I_{n_k = n \wedge m_n = m} B_m^n dP \\ &= \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} \int \left(\int I_{n_k = n \wedge m_n = m} (x_1, \dots, x_n; z_{n+1}, \dots) B_m^n (x_1, \dots, x_n; z_{n+1}, \dots) dP^n(x_1, \dots, x_n) \right) dP_n. \end{aligned}$$

By Lemma 6, $I_{n_k = n \wedge m_n = m}$ is $\mathcal{F}_n$ -measurable and thus constant with respect to the arguments

$z_{n+1}, z_{n+2}, \dots$. Thus, we continue

$$= \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} \int I_{n_k=n \wedge m_n=m}(x_1, \dots, x_n) \left(\int B_m^n(x_1, \dots, x_n; z_{n+1}, \dots) dP^n(x_1, \dots, x_n) \right) dP_n.$$

Now $I_{n_k=n \wedge m_n=m}(x) = 1$ implies a fortiori that there is a bet along x at time n on F_m^n and so

$\int B_m^n(x_1, \dots, x_n; z_{n+1}, \dots) dP^n(x_1, \dots, x_n) = 1$. Thus, we continue

$$\begin{aligned} &= \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} \int I_{n_k=n \wedge m_n=m}(x_1, \dots, x_n) dP_n = \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} P[n_k = n \text{ and } m_n = m] \\ &= \sum_{n=1}^{\infty} P[n_k = n \text{ and } m_n < \infty] = \sum_{n=1}^{\infty} P[n_k = n] = P[n_k < \infty]. \end{aligned}$$

iii): We show the result for $j = k + 1$. The same argument applies for any $j > k$. First we observe that

$$\begin{aligned} \text{cov}(B_k^P, B_{k+1}^P) &= \int B_k^P(B_{k+1}^P - 1) dP - \int (B_{k+1}^P - 1) dP \\ &= \int B_k^P(B_{k+1}^P - 1) dP = \sum_{\bar{n}=1}^{\infty} \int I_{n_{k+1}=\bar{n}} B_k(B_{k+1} - 1) dP. \end{aligned}$$

Hence, it suffices to show $\int I_{n_{k+1}=\bar{n}} B_k(B_{k+1} - 1) dP = 0$ for arbitrary $\bar{n}$.

The first step is to show that if we restrict attention to x whose $k + 1^{\text{th}}$ bet is at $\bar{n}$, then the k^{th} bet is a $\mathcal{F}_{\bar{n}}$ -measurable function. More precisely, $I_{n_{k+1}=\bar{n}} B_k$ is a $\mathcal{F}_{\bar{n}}$ -measurable function since

$$I_{n_{k+1}=\bar{n}} B_k = I_{n_{k+1}=\bar{n}} \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} I_{n_k=n \wedge m_n=m} B_m^n = I_{n_{k+1}=\bar{n}} \sum_{n=1}^{\bar{n}} \sum_{m: r_m^n \leq \bar{n}} I_{n_k=n \wedge m_n=m} B_m^n$$

and $\forall m, n$, s.t. $n, r_m^n < \bar{n}$, $\{n_{k+1} = \bar{n}\}$, $\{n_k = n\}$ and $\{m_n = m\}$ are all $\mathcal{F}_{\bar{n}}$ -measurable sets and

by Lemma 4, B_m^n is a $\mathcal{F}_{\bar{n}}$ -measurable function.

The fact that $I_{n_{k+1}=\bar{n}}B_k$ is a $\mathcal{B}_{\bar{n}}$ -measurable function allows us to “pull out” B_k from the inside integral when parsing across $\bar{n}$ length histories:

$$\begin{aligned} & \int I_{n_{k+1}=\bar{n}}B_k(B_{k+1}-1) dP \\ &= \int I_{n_{k+1}=\bar{n}}(x_1, \dots, x_{\bar{n}})B_k(x_1, \dots, x_{\bar{n}})\left(\int B_{k+1}(x_1, \dots, x_{\bar{n}}, z_{\bar{n}+1}, \dots) dP^{\bar{n}}(x_1, \dots, x_{\bar{n}})-1\right) dP_{\bar{n}} = 0. \end{aligned}$$

To see that the last line is zero, note that when $n_{k+1}(x) = \bar{n}$, there is a bet at time $\bar{n}$ along x on the $m_{\bar{n}}$ event and so

$$\int B_{k+1}(x_1, \dots, x_{\bar{n}}, z_{\bar{n}+1}, \dots) dP^{\bar{n}}(x_1, \dots, x_{\bar{n}}) = \int B_{m_{\bar{n}}}^{\bar{n}}(x_1, \dots, x_{\bar{n}}, z_{\bar{n}+1}, \dots) dP^{\bar{n}}(x_1, \dots, x_{\bar{n}}) = 1 \text{ by Lemma 4.}$$

The last step is to apply the law of large numbers for uncorrelated random variables as in the main text to obtain the result in Theorem 3. ■

9. REFERENCES

- Ash, R. (1972), *Real Analysis and Probability*, Boston, Academic Press, Inc.
- Billingsley, P. (1995), *Probability and Measure* (3rd ed.), New York, John Wiley & Sons.
- Blackwell, D and L. Dubins (1962), "Merging of Opinions with Increasing Information," *Annals of Mathematical Statistics*, 38:882-886.
- Kalai, E. and E. Lehrer (1993), "Rational Learning Leads to Nash Equilibrium," *Econometrica*, 61: 1019-1045.
- Nachbar, J. (1995), "Prediction, Optimization, and Rational Learning in Games," mimeo, Department of Economics, Washington University.
- Shiryaev, A. N. (1996), *Probability (Second Edition)*, New York, Springer, Chapter VII.
- Stokey, N. and R. E. Lucas Jr. (1989), *Recursive Methods in Economic Dynamics*, Cambridge, Harvard University Press.

10. NOTES

* We would like to thank Kyle Bagwell, Subir Bose, Phoebus Dhrymes, Prajit Dutta, Ehud Lehrer, Andrew Newman, and the associate editor and referees of a previous version.

¹ Blackwell and Dubins' restrict attention to probability measures that are fully described by a collection of nodal measures (See the appendix). Conversely, every collection of nodal measures produces a (unique) measure on full paths (See, e.g. Ash (1972), p. 109).

² See the appendix for the general case.

³ Here and in the next sentence, we abuse notation in the usual manner by viewing h_n as both an element of the partial histories X_n and a subset of the full histories X .

⁴ See the appendix for the general case.

⁵ See the appendix for the general case.

⁶ This is shown for the more general case in the appendix.

⁷ An alternative approach here is to: 1) recognize that, with appropriate choice of filtration, B_k is a submartingale difference sequence for P and a supermartingale difference sequence for Q , 2) consider the Doob decompositions (See, e.g. Shiryaev, p. 482) of the corresponding sub- and supermartingales, and 3) apply the martingale law of large numbers (e.g. Shiryaev, p. 501) to the martingale components of these decompositions. Note that this "martingale approach" to our proof does *not* call upon Doob's martingale *convergence* result, as do the proofs in Blackwell and Dubins and Kalai and Lehrer. In particular the martingale law of large numbers follows from martingale inequalities.

⁸ There seems to be an error in the fourth line from the end of the proof in this text, which does not affect the truth of the proposition. Apparently, it should read:

$$“\mu(Z_k \geq k^2 \varepsilon) \leq 4M / k^2 \varepsilon^2, \text{ all } \varepsilon > 0, \text{ all } k.”$$

⁹ While the “Grain of Truth” assumption only makes sense following restriction to a countable set of possible measures, there exist certain kinds of restrictions to uncountable sets in which the more general absolute continuity assumption is reasonable. The classic example, as pointed out by Blackwell and Dubins, is restriction to the class of measures that are i.i.d. with a finite number of parameters.

¹⁰ That is to say, $\mathcal{B}(i)$ is the smallest σ -algebra containing some countable collection of subsets of $X(i)$.

¹¹ Indeed, the value of the portion of Blackwell and Dubins’ theorem that is not contained in our own result is unclear: the fact there exists a version of Q^n that converges to a given version of P^n is difficult to interpret when each $\mathcal{B}(i)$ is *not* countably generated. For when the $\mathcal{B}(i)$ are not countably generated, two different pairs of choices P^n, Q^n and $\hat{P}^n, \hat{Q}^n$ may yield distance functions $|P^n(x_1, \dots, x_n) - Q^n(x_1, \dots, x_n)|$ and $|\hat{P}^n(x_1, \dots, x_n) - \hat{Q}^n(x_1, \dots, x_n)|$ that differ on a set of both P and Q positive probability. In other words, without countable generation, the distance norm and hence the convergence properties Q and P will not be invariant to our choice of P^n and Q^n . The question then arises: if given P^n , some choices of Q^n converge to P^n and some do not, has there been a “merging of opinions?”

¹² For the discrete σ -algebra and so for all σ -algebras on $N \cup \{\infty\}$.

1996-1997 Discussion Paper Series

Department of Economics
Columbia University
1022 International Affairs Bldg.
420 West 118th Street
New York, N.Y., 10027

The following papers are published in the 1996-97 Columbia University Discussion Paper series which runs from early November to October 31 of the following year (Academic Year).

Website: <http://www.columbia.edu/dlc/wp/econ/index.html>.

You may download any papers found on this site.

For Ordering Hardcopies:

Domestic orders for discussion papers are available for purchase at the cost of \$8.00 (U.S.) Per paper and \$140.00 (US) for the series.

Foreign orders cost \$10.00 (US) per paper and \$185.00 for the series.

To order discussion papers, please write to the Discussion Paper Coordinator at the above address along with a check for the appropriate amount, made payable to Department of Economics, Columbia University. Please be sure to include the series number of the requested paper when you place an order.

1996-97 Discussion Paper Series

9697-01	Fertility Behavior Under Income Uncertainty	by: P. Ranjan
9697-02	Trade Restrictions, imperfect Competition and National Welfare with Foreign Capital Inflows	by: P. Ranjan
9697-03	Restructuring an Industry during Transition: a Two-Priced Model	by: R. Ericson
9697-04	A Conformity Test for Cointegration.	by: P. Dhrymes
9697-05	Low-Wage Employment Subsidies in a Labor-Turnover Model of the 'Natural Rate' (November 1996)	by: H.T. Hoon E. Phelps
9697-06	The Knowledge Revolution	by: G. Chichilnisky
9697-07	The Role of Absolute Continuity in "Merging Opinions" and "Rational Learning"	by: R. Miller C.W. Sanchirico