

UC Berkeley

UC Berkeley Previously Published Works

Title

Using Prediction Market Data to Illustrate Undergraduate Probability

Permalink

<https://escholarship.org/uc/item/0s22057q>

Journal

American Mathematical Monthly, 120(7)

ISSN

0002-9890

Author

Aldous, David J

Publication Date

2013

DOI

10.4169/amer.math.monthly.120.07.583

Peer reviewed

Using Prediction Market Data to Illustrate Undergraduate Probability

David J. Aldous

Abstract. Prediction markets provide a rare setting, where results of mathematical probability theory can be related to events of real-world interest and where theory can be compared to data. The paper discusses two simple mathematical results—the halftime price principle and the serious candidates principle—and corresponding data from baseball and the 2012 Republican Presidential nomination race.

1. INTRODUCTION. Outside the context of games of chance based on artifacts with physical symmetry (dice, lotteries, roulette wheels, etc.), it is surprisingly difficult to find interesting real-world data that illustrates undergraduate probability calculations. The majority of examples and exercises in undergraduate probability textbooks are either “just mathematics” (X ’s and Y ’s without attempting a real-world story) or unmotivated (“consider an urn containing 5 black balls, 3 white balls, and 2 blue balls”) or enter a curious fantasy world. In this fantasy world, tossed coins can be biased, darts can hit a board, or a stick can be split with uniform distribution, and people choose majors or desserts at random. But none of these are any more realistic than Jedi Knights or spherical cows. A notable exception is the recent book by Grinstead, Peterson, and Snell [3], which gives a careful introduction to four topics (streaks, the stock market, lotteries, and fingerprints), combining mathematical probability models with relevant data. The purpose of this article is to publicize another topic, prediction markets. We state below two very simple results from mathematical theory, and will give their mathematical derivation and compare with some data. As well as being (hopefully) interesting to MONTHLY readers, we might regard this article (in the spirit of [3]) as supplementary material for an undergraduate probability course.

A prediction market (browse <http://www.intrade.com> for better understanding) is essentially a venue for betting whether a specified event will occur (perhaps before a specified time), where the betting is conducted via participants buying and selling contracts with each other rather than with the operators of the market. In other words, it is structured like a stock market rather than a bookmaker. As we will discuss briefly in section 4, the mathematics of prediction markets is very similar to that of stock markets, but in several respects prediction markets are conceptually simpler and therefore more suitable for an introductory-level treatment. In a typical prediction market (details vary between markets in unimportant ways) a contract will expire at 100 or 0, depending on whether or not the specified event occurs. At any time you can see posted offers to buy and to sell, say at 56.8 and 57.2, and if you wish to buy a contract you can either pay the posted 57.2 and buy instantly, or post an offer to pay (say) 57.0 and wait to find out if anyone is willing to sell at that price. The key conceptual point is that the current market price, 57 in this example, can be interpreted as a consensus probability (57 percent, of course) that the event will happen—see section 3.2 for discussion. The special feature of prediction markets, as a topic within undergraduate mathematical probability, is as the essentially unique context in which

<http://dx.doi.org/10.4169/amer.math.monthly.120.07.583>
MSC: Primary 60G44

we have substantial real-world data interpretable as probabilities and fluctuations of probabilities over time. Here we refer to the “price” of a contract, which the reader must always interpret as “the probability of the event occurring, given what is currently known”. The price, obviously, will fluctuate as relevant real-world information becomes known; how it will fluctuate is unknown in advance, and hence random. We use “price” rather than “probability” to avoid the linguistic confusion involved when talking about probabilities of probabilities.

Figure 1 shows price fluctuations of two specific contracts. These involved two different contexts and time scales for prediction market contracts—a few hours’ duration for a sports match or a year-long run-up to an election. The first concerned which team would be the winner of a baseball game (New York Mets vs. Florida Marlins) on August 31, 2008; the second concerned whether Newt Gingrich would be chosen as the Republican U.S. Presidential nominee in 2012. In the charts, the horizontal scale is time, and the vertical scale (on the right) is price (note that in neither case is the whole range [0, 100] shown). The low bars in the second chart show daily volume, typically under a thousand \$10 contracts per day. See <http://www.intrade.com> for clear examples of current charts.

We now state two “principles”, by which we mean assertions based solely on mathematical arguments, about how prices in prediction markets should behave.



Figure 1. Two prediction market price charts

The halftime price principle. In a sports match between equally good teams, at half-time there is some (prediction market) price for the home team winning. This price varies from match to match, depending largely on the scoring in the first half of the match. Theory says that its distribution should be approximately uniform on $[0, 100]$.

The serious candidates principle. Consider an upcoming election with several candidates, and a (prediction market) price for each candidate. Suppose initially that all these prices are below b , for given $0 < b < 100$. Theory says that the expected number of candidates whose price ever exceeds b equals $100/b$.

The mathematics involved here is simple and has undoubtedly been folklore for generations, though we are not aware of previous discussion in the spirit of this article. We invented the names above: Candidates become *serious* (rather than fringe) when their chance of election exceeds some threshold.

For each principle we will explain the theory and then show some data. We assume that the reader is familiar with basic notions from an undergraduate textbook such as [4, 7, 8]. Explaining the second principle will involve the concept of a *martingale*, typically not encountered until a second course in probability (e.g., Chapter 6 of [9] or Chapter 6 of [5]), but ultimately so widely useful that a graduate course on mathematical probability can be taught with martingales as the central topic [10]. We give only the minimal account of martingales needed here. Every textbook introduction to martingales we know treats them as mathematical objects without explicit relation to real-world data, but an instructor of a course introducing martingales could continue our style of illustrating other mathematical results about martingales with real-world prediction market data, and some examples can be found in [1].

2. THE HALFTIME PRICE PRINCIPLE. To elaborate this principle we imagine a sport in which (like almost all team sports) the result is decided by point difference, and for simplicity, imagine a sport like baseball or American football where there is a definite winner (ties are impossible or rare). Also for simplicity, we assume that the teams are equally good, in the sense that there is initially a 50% probability of the home team winning (that is, equally good after taking home-field advantage into account). Write Z_1 for the point difference (points scored by home team, minus points scored by visiting team) in the first half, and Z_2 for the point difference in the second half.

A fairly realistic mathematical model of this scenario is to assume that:

- (i) Z_1 and Z_2 are independent random variables, with the same distribution;
- (ii) their distribution is symmetric about zero, that is, their distribution function $F(z)$ satisfies $F(z) = 1 - F(-z)$. For mathematical ease we add an unrealistic assumption (to be discussed later);
- (iii) the distribution is continuous.

Under these assumptions we can do a calculation, though first we recall the slightly sophisticated notation that treats conditional probabilities as random variables. For an event A and a random variable Y , the elementary notation for conditional probabilities

$$P(A | Y = y) = g(y) \text{ for all } y$$

(the left side is always *some* function of y) can be rewritten as

$$P(A | Y) = g(Y). \tag{1}$$

The following calculation exemplifies the usefulness of this notation.

The probability that the home team wins, given that the first-half point difference is z , is

$$\begin{aligned}
 P(Z_1 + Z_2 > 0 \mid Z_1 = z) &= P(Z_2 > -z) \text{ by independence} \\
 &= F(z) \text{ by symmetry}
 \end{aligned}$$

and therefore the price at halftime, which is the conditional probability of the home team winning, given the observed value of Z_1 , is

$$P(Z_1 + Z_2 > 0 \mid Z_1) = F(Z_1). \tag{2}$$

But in fact (e.g., [8, p. 234]), for a continuous distribution it is always true that $F(Z_1)$ has uniform distribution on (0%, 100%).

That is the mathematical justification for the principle. We can think of various defects in the model, most obviously the fact that in real sports the points are integer-valued, but for reasons explained below, we suspect that this does not make a huge difference, even in the worst case of a low-scoring sport like soccer.

2.1. A little data.

Errors using inadequate data are much less than those using no data at all. [Charles Babbage]

In the baseball match chart in Figure 1, the initial price was near 50; the price at halftime (for baseball we simply used halfway through the match duration) was around 62.

In 30 baseball games from 2008 for which we have the prediction market prices as in Figure 1, and for which the initial price was around 50%, the prices (as percentages) halfway through the match were as follows:

- 07, 10, 12, 16, 23, 27, 31, 32, 33, 35, 36, 38, 40, 44, 46
- 50, 55, 57, 62, 65, 70, 70, 71, 73, 74, 74, 76, 79, 89, 93.

Figure 2 (left) compares the distribution function of this data to the (straight line) distribution function of the uniform distribution.




Figure 2. The left diagram shows the empirical distribution function for the baseball data, compared with the uniform distribution. The right diagram shows the theoretical distribution function in the soccer model, again compared with the uniform distribution.

The data appears roughly consistent with our halftime price principle. We do not attempt formal tests of significance (“goodness of fit”), which are not informative in our context of an approximate theoretical prediction and limited data.

Caveat. The simplicity of the stated halftime price principle depends on the teams being equally good. For unequal teams, the distribution of halftime price will depend on the distribution of the point differences Z_i as well as the initial price.

2.2. A soccer model. To investigate the effect of discrete points theoretically, take a standard model for soccer, where we suppose that the teams score goals at the times of independent Poisson processes of rates λ_1 and λ_2 per match-duration, with a shoot-out to determine the winner if the score is tied at the end. We can readily adapt the previous calculation to this model. Taking, for instance, $\lambda_1 = \lambda_2 = 2$, Figure 2 (right) shows the distribution function of the halftime price in this model, compared to the (straight line) distribution function of the uniform distribution.

Here, the discrete distribution of the halftime price arises from the discrete distribution of the halftime point difference (likely to be 1 or 0 or -1). This reminds us of other defects of the model. In practice, the prediction market prices depend not only on halftime score, but also on other factors, such as quantitative (e.g., shots on goal) and qualitative assessments of each team’s play in the first half. These vary from match to match and will tend to smooth out the distribution. We conjecture that data on halftime soccer prices for equally-matched teams would, in fact, have a roughly uniform distribution, as with the baseball data above.

3. PREDICTION MARKETS AND MARTINGALES. From the very broad field of martingale theory, let us extract several points to emphasize.

1. The notion of your successive *fortunes* (amounts of money you have) during a sequence of bets at fair odds can be formalized mathematically as a *martingale*. The gambling interpretation enables proofs of theorems concerning martingales to be expressed in very intuitive language. Then the mathematical definition and theorems can be used (if their hypotheses are satisfied) for random processes arising in contexts completely unrelated to money or gambling.
2. One theorem about martingales says that the overall result of any system for deciding how much and when to bet, within this “fair odds” setting, is simply equivalent to a single bet at fair odds. So we can prove theorems about martingales by inventing hypothetical betting systems and analyzing their possible outcomes.
3. There are plausible reasons to believe that prediction market prices should behave like martingales.

In the next two sections we say a few words about these points. The reader willing to accept them may jump ahead to section 3.4, where we use them to derive the *serious candidates principle*.

3.1. Martingales. For our purposes, a *fair bet* (more accurately, a bet at fair odds) is one in which the expectation of your monetary gain G equals zero; that is,

$$E[G] = 0,$$

where a loss is a negative gain. This ignores issues of utility and risk-aversion, which we won’t consider. In other words, in order for you to receive from me a random payoff

X in the near future, the “fair” amount you should pay me now is $E[X]$, because then your gain (and my loss) $X - E[X]$ has expectation zero. If a bet is fair, then doubling the stake and payoff, or multiplying both by -3 to bet in the opposite direction, is again a fair bet.

A formal definition of *martingale* is a process, that is, a sequence of real-valued random variables, satisfying for each $n \geq 0$

$$E(X_{n+1} \mid X_n = x_n, X_{n-1} = x_{n-1}, \dots, X_0 = x_0) = x_n, \text{ for all } x_0, x_1, \dots, x_n. \quad (3)$$

This is pretty hard to interpret if you’re not familiar with the probability notation, so we’ll try to explain in words, in the context of gambling. Imagine a person making a sequence of bets, and after the n th bet is settled, his fortune is x_n . After placing the next bet but before knowing the outcome, the gain G_{n+1} on that bet is random, and (3) says that

$$E(G_{n+1} \mid X_n = x_n, X_{n-1} = x_{n-1}, \dots, X_0 = x_0) = 0,$$

i.e., that the expected gain on the bet, given what we currently know, equals zero—the “fair” concept.

A textbook example ([6, ex. 10.2.6]) of a martingale (X_n) arising in a context unrelated to money or gambling, concerns the Wright–Fisher model in population genetics. The model shows (without mutation or selection) that the proportion X_n of genes in generation n , which are a particular allele, forms a martingale.

Developing the basic mathematics of martingales requires many small steps to introduce and explain notation. Below we give a verbal overview and refer the reader to the advanced textbook [10] for the mathematics.

Return to the gambling story above, where another gambler’s fortune is the martingale (3). Imagine that you are copying or modifying the bets of this other gambler. A simple way to do so is to copy exactly what the gambler does, but stop after the T th bet is resolved, where T can be chosen on the fly, that is, depending on what has happened so far, but not foreseeing the future. It is perhaps remarkable that there is a precise mathematical definition ([10, 10.8]) of a *stopping time* T capturing this idea. Following this system, your gain is $X_T - x_0$. The basic form of the *optional sampling theorem* for martingales ([10, A14]) says that

$$E[X_T] = x_0 \text{ for each stopping time } T.$$

In the gambling context, this says that despite the fact that you are using a “system” (in this case just some rule for when to stop), your net result is a fair bet. (This theorem and the theorem below have side conditions that are automatically satisfied in our settings.)

As a very general way of copying another gambler, on the n th bet (for each n) stake some multiple H_n of the other gambler’s stake, where H_n may depend on the past, but cannot foresee the result of the n th bet. Following such a system, your gain Y_n is determined by the processes (X_n) and (H_n) via the formula $Y_{n+1} - Y_n = H_n(X_{n+1} - X_n)$ and is called a *martingale transform* ([10, 10.6]) or *discrete stochastic integral*. The key fact is that Y_n behaves as a martingale, and that whenever you choose to stop, your gain Y_T has expectation zero ([10, 10.7]). The latter result is often referred to via a phrase like “impossibility of gambling systems,” but we would prefer a more positive and informative name, so let us follow [2] and call it the **conservation of fairness theorem**.

3.2. Prediction market prices are approximately martingales. Definitions and theorems about martingales, as outlined above, can be regarded as a part of pure mathematics, with the references to gambling being just a side story to aid intuition. To now argue that *there are plausible reasons to believe that prediction market prices should behave like martingales*, we must obviously leave pure mathematics at some point. Indeed, any serious treatment would enter realms of philosophy, psychology, economics, and empirical data. Here, we focus on where exactly pure mathematics ends and other issues begin.

First, recall that general mathematical results about probabilities and conditional probabilities of *events* can be derived from those for expectations and conditional expectations of *random variables*, by the device of identifying an event A with its $\{0, 1\}$ -valued indicator random variable 1_A . Second, outside very simple settings, probabilities depend on “information known at the current time n ”, and the formalization of this notion within the usual axioms of mathematical probability is as a collection (a *sigma-algebra* or *sigma-field*, technically) of events, conventionally denoted by \mathcal{F}_n , whose outcomes we know. For such a collection \mathcal{F} and any event A , we can define the conditional probability $P(A | \mathcal{F})$ as a random variable, extending the notation (1) where the “information” in \mathcal{F} is the value of Y . Here $P(A | \mathcal{F})$ is random in the “prior” sense—before we know which events in \mathcal{F} actually happened.

A benefit of going through this abstract setup is an easy theorem ([10, 10.4c]) which says that, for any event A and any sequence \mathcal{F}_n representing “information known at the current time n ,” the conditional probabilities $X_n := P(A | \mathcal{F}_n)$ always form a martingale.

This is about as far as pure mathematics can take us. Concerning probabilities for the kind of interesting future real-world events exemplified by results of sports matches or elections, there is longstanding philosophical debate about whether such probabilities can or should be interpreted as anything other than subjective opinions. But perhaps a more substantial issue is that the way you or I might assess probabilities for such events, though based on something that might be called “information”, is manifestly not the way envisaged in the axiomatic setup; this would involve first setting out all the possible relevant events that (from the standpoint of some past time) might have happened by now or in the future, then assigning probabilities to every combination of events happening and not happening, then looking at which of these events did or did not happen by now, and finally doing the required calculation.

In a real prediction market, different individual participants will assess probabilities somewhat differently, and (among those willing to bet actual money) the market price represents a balance point between willing buyers and willing sellers; we can call this price a “consensus probability”. So the central issue is: *Why should such consensus probabilities change in time in the same way as conditional probabilities, within the axiomatic setup of mathematical probability?* Typical verbal arguments use an undefined notion of “information” and simply jump over this issue, and we don’t know of any satisfactory argument. So it seems most appropriate to call the assertion that *prediction market prices should behave like martingales*, a *hypothesis*, and seek to see if its mathematical consequences are consistent with empirical data. Obviously, this is similar to the *efficient market hypothesis* in finance, though as discussed in section 4 the setting of prediction markets is conceptually simpler than stock markets.

3.3. Were there improbably many candidates for the 2012 Republican nomination whose fortunes rose and fell? In the race for the 2012 Republican Presidential nomination there were many candidates whose popularity rose and then fell noticeably—Donald Trump, Newt Gingrich, Sarah Palin, Rick Perry, and Michelle

Bachmann, for instance. Almost all discussions of the race have shared the presumption that the number of such candidates was much larger than usual, and speculated on the reasons, e.g., an “anyone but Romney” sentiment. But is that presumption true?

We need to distinguish between two meanings. Opinion polls ask questions such as, “if you were voting tomorrow, who would you vote for?”. Mathematics says nothing about how much such opinions may fluctuate over a year-long campaign, just as mathematics says nothing about how much fashions in popular music may fluctuate. We could devise some statistic to measure these fluctuations and compare it empirically with the statistics from previous races, but we cannot compare it to any theoretical prediction.

On the other hand, the theoretical argument that every prediction market price should be a martingale is not affected by fashion or opinion poll results. So we can examine whether the prediction market prices in this particular race behaved differently from how theory says prediction market prices should behave, which would be an indication of some unusual aspect of the 2012 race.

3.4. Argument for the serious candidates principle. We want a model for prediction market prices for an upcoming election, in the generalized sense of one candidate being chosen at a specified future date (so this covers future Oscar winners, for instance, for which Intrade also provides markets). The only assumption we need is that each candidate’s price is a continuous-path martingale. Here, continuous-path is not literally true (prices are discrete) but corresponds to the idea of a “liquid market” with small spread between bid and ask prices, which is reasonably accurate for the election markets under consideration.

To restate the serious candidates principle:

Consider an upcoming election with several candidates, and a (prediction market) price for each candidate. Suppose initially that all these prices are below b , for given $0 < b < 100$. Theory says that the expected number of candidates whose price ever exceeds b , equals $100/b$.

Here is the mathematical argument, based on a hypothetical betting system. For each candidate, buy a contract on that candidate if and when their price reaches b . The total cost of these contracts is bN_b , where N_b is the random number of candidates whose price ever reaches level b . Exactly one candidate is elected, and your contract on that candidate earns you 100. So your gain is $100 - bN_b$. The conservation of fairness theorem says that the expected gain equals zero, and that the equation $E[100 - bN_b] = 0$ rearranges to $E[N_b] = 100/b$.

3.5. Data. Table 1 shows the maximum Intrade prediction market price (up to June 8) for each of the 16 leading candidates for the 2012 Republican Presidential nomination.

Table 1. Maximum prediction market prices.

Romney	98	Perry	39	Gingrich	38	Palin	28
Pawlenty	25	Santorum	18	Huntsman	18	Bachmann	18
Huckabee	17	Daniels	14	Christie	10	Giuliani	10
Bush	9	Cain	9	Trump	8.7	Paul	8.5

These numbers might well suggest to a non-mathematician that the number of sometime-serious candidates was unusually large. But look at Table 2, which compares observed data with the mathematical prediction for “number of candidates with maximum price $\geq b$ ” for several values of b .

Table 2 indicates that the number of candidates whose fortunes rose and fell in this “probability of winning” sense was scarcely more than would be expected on mathematical grounds.

Table 2. Observed and expected numbers exceeding threshold prices.

	Expected	Observed
$b = 33.3$	3	3
$b = 20$	5	5
$b = 16.6$	6	9
$b = 12.5$	8	10

Two technical points. In Table 2 we used $100/b$ as “expected”, without considering whether some initial prices might have been greater than b . Data on initial prices is somewhat unreliable (because the contract may initially be thinly traded), but the only candidate whose initial price was clearly above 10 was Romney, at about 23. Correcting for this would make the “expected” numbers slightly smaller for small b . Of course, for a campaign where two candidates started with price 40, the “expected” numbers would be very different. Another important general point is that, for long-duration contracts, low prediction market prices overstate the true consensus probability because of the “covering your position” requirement. That is, even if you were certain an event would not happen, you might not be willing to sell a contract for 3 because your sure gain of 3 is offset by the opportunity or interest cost of the market requirement that you deposit 97 to cover a possible loss. Correcting for this effect would make the “expected” numbers in Table 2 larger than shown for small b .

A bottom-line conclusion. To the extent that mathematics can say anything relevant, it says that the fundamental driving feature of the 2012 nominee campaign was that it started without any clear favorite. The subsequent fluctuations were then consistent with what theory predicts. In other words, even if it is actually true that the month-to-month fluctuations in opinion poll standings were greater than usual, we can see no sign that this unduly influenced the smart money being wagered on the prediction market.

3.6. Another check of theory and data. A mathematician familiar with martingale theory might look at the Figure 1 chart for Newt Gingrich and wonder if it shows too many fluctuations to be plausibly a martingale. For instance, the chart shows two separate downcrossings from 20 to 10, in December 2011 and in late January 2012. This mathematician has in mind the *upcrossing inequality* ([10, 11.3]), which limits the likely number of such crossings. We can conduct another check of theory versus data by considering crossings. The relevant theory turns out to be:

Consider a price interval $0 < a < b < 100$. Then consider an upcoming election with several candidates, and a (prediction market) price for each candidate,

where initially all these prices are below b . Theory says that the expected total number of downcrossings of prices (sum the numbers for each candidate) over the interval $[a, b]$ equals $(100 - b)/(b - a)$.

To gather data for the interval $[10, 20]$, we need only look at the five candidates in Table 1 whose maximum price exceeded 20. Their numbers of downcrossings of $[10, 20]$ were:

Palin (2), Romney (0), Perry (1), Pawlenty (2), Gingrich (2).

So the observed total 7 is in fact close to the theoretical expectation of 8. To derive the formula quoted, we again consider a hypothetical betting system. For each candidate, buy a contract on that candidate if and when their price reaches b . If the price subsequently falls to a , then sell; but buy again if the price reaches b , and continue. Exactly one of these contracts will expire at 100, and the others will be sold at price a , the number $D_{a,b}$ of these others being the number of downcrossings of $[a, b]$. So your gain is $(100 - b) - D_{a,b}(b - a)$. The conservation of fairness theorem says the expected gain equals zero, and the equation $E[(100 - b) - D_{a,b}(b - a)] = 0$ rearranges to $E[D_{a,b}] = (100 - b)/(b - a)$.

4. COMPARING PREDICTION MARKETS AND STOCK MARKETS. We asserted that prediction markets are conceptually simpler than stock markets, so let us finish by making some comparisons between the two.

1. In both markets, the “market price” is by definition the price at which buyers and sellers are willing to trade. Assigning any other interpretation to the price of one share of Apple corporation is a matter of debate—one interpretation from the rationalist school would be that it represents a consensus estimate of discounted future earnings, adjusted by an equity risk premium whose size depends on the risk premiums imputed to alternative investments. In contrast, the interpretation of a prediction market price as the probability of the specified event is much more definite.
2. The price in a prediction market must be between 0 and 100, and will expire at 0 or 100 at a known time determined by an explicit event outside the market.
3. A prediction market is mathematically simpler because we need no empirical data to make the theoretical predictions; for the analogous predictions in a stock market, we need an estimate of variance rate.
4. Compared to stock markets, prediction markets are often thinly traded, suggesting that they will be less efficient and less martingale-like.
5. Standard economic theory asserts that long-term gains in a stock market will exceed long-term rewards in a non-risky investment, because investors’ risk-taking must be rewarded. In this picture, a stock market is a “positive sum game” benefiting both investors and corporations seeking capital; financial intermediaries and speculators earn their share of the gain by providing liquidity and convenient diversification for investors. In contrast, the prediction markets currently in operation are too small to have substantial effect on the real economy, and so are zero-sum, in fact, slightly negative-sum because of transaction costs.

ACKNOWLEDGMENTS. Thanks to Mykhaylo Shkolnikov and Andrew Gelman for comments.

REFERENCES

1. D. J. Aldous, On Chance and Unpredictability: 20 lectures on the links between mathematical probability and the real world, 2012, draft at <http://www.stat.berkeley.edu/~aldous>.
2. S. N. Ethier, *The Doctrine of Chances: Probabilistic Aspects of Gambling*, Springer-Verlag, Berlin, 2010.
3. C. M. Grinstead, W. P. Peterson, J. L. Snell, *Probability Tales*, American Mathematical Society, Providence, RI, 2011.
4. C. M. Grinstead and J. L. Snell, *Introduction to Probability*, 2nd ed., American Mathematical Society, Providence, RI, 1997.
5. S. Karlin, H. M. Taylor, *A First Course in Stochastic Processes*, 2nd ed., Academic Press, New York, 1975.
6. K. Lange, *Applied Probability*, 2nd ed., Springer, New York, 2010.
7. J. Pitman, *Probability*, Springer, New York, 1993.
8. S. Ross, *A First Course in Probability*, 6th ed., Prentice Hall, Upper Saddle River, NJ, 2002.
9. S. Ross, *Stochastic Processes*, 2nd ed., Wiley, New York, 1996.
10. D. Williams, *Probability with Martingales*, Cambridge University Press, Cambridge, 1991.

DAVID ALDOUS received his Ph.D. from the University of Cambridge in 1977, and since 1979 has been at U.C. Berkeley. After a career doing technical research within mathematical probability, he is now interested in relating mathematics to real-world data. Outside of work he drinks lattes, reads science fiction and *The Economist*, drinks wine, walks the dog, and wastes too much time playing *Civilization III*.
Department of Statistics, 367 Evans Hall, U.C. Berkeley CA 94720-3860
aldous@stat.berkeley.edu

“I see some parallels between the shifts of fashion in mathematics and in music. In music, the popular new styles of jazz and rock became fashionable a little earlier than the new mathematical styles of chaos and complexity theory. Jazz and rock were long despised by classical musicians, but have emerged as art-forms more accessible than classical music to a wide section of the public. Jazz and rock are no longer to be despised as passing fads. Neither are chaos and complexity theory. But still, classical music and classical mathematics are not dead. Mozart lives, and so does Euler. When the wheel of fashion turns once more, quantum mechanics and hard analysis will once again be in style.”

From Freeman Dyson’s review of *Nature’s Numbers*
 by Ian Stewart (Basic Books, 1995).
 (From p. 612 in the *American Mathematical Monthly*,
 August–September 1996)

—Submitted by Jon Borwein