

Randomness and Coincidences: Reconciling Intuition and Probability Theory

Thomas L. Griffiths & Joshua B. Tenenbaum
Department of Psychology
Stanford University
Stanford, CA 94305-2130 USA
{gruffydd, jbt}@psych.stanford.edu

Abstract

We argue that the apparent inconsistency between people's intuitions about chance and the normative predictions of probability theory, as expressed in judgments about randomness and coincidences, can be resolved by focussing on the evidence observations provide about the processes that generated them rather than their likelihood. This argument is supported by probabilistic modeling of sequence and number production, together with two experiments that examine judgments about coincidences.

People are notoriously inaccurate in their judgments about randomness, such as whether a sequence of heads and tails like HHTHTTTH is more random than the sequence HHHHHHHH. Intuitively, the former sequence seems more random, but both sequences are equally likely to be produced by a random generating process that chooses H or T with equal probability, such as a fair coin. This kind of question is often used to illustrate how our intuitions about chance deviate from the normative standards set by probability theory. Our intuitions about coincidental events, which seem to be defined by their improbability, have faced similar criticism from statisticians (eg. Diaconis & Mosteller, 1989).

The apparent inconsistency between our intuitions about chance and the formal structure of probability theory has provoked attention from philosophers and mathematicians, as well as psychologists. As a result, a number of definitions of randomness exist in both the mathematical (eg. Chaitin, 2001; Kac, 1983; Li & Vitanyi, 1997) and the psychological (eg. Falk, 1981; Lopes, 1982) literature. These definitions vary in how well they satisfy our intuitions, and can be hard to reconcile with probability theory. In this paper, we will argue that there is a natural relationship between people's intuitions about chance and the normative standards of probability theory. Traditional criticism of people's intuitions about chance has focused on the fact that people are poor estimators of the likelihood of events being produced by a particular generating process. The models we present turn this question around, asking how much more likely a set of events makes a particular generating process. This question may be far more useful in natural inference situations, where it is often more important to reason diagnostically than predictively, attempting to infer the structure of our world from the data we observe.

Randomness

Reichenbach (1934/1949) is credited with having first suggested that mathematical novices will be unable to produce random sequences, instead showing a tendency to overestimate the frequency with which outcomes alternate. Subsequent research has provided support for this claim (reviewed in Bar-Hillel & Wagenaar, 1991; Tune, 1964; Wagenaar, 1972), with both sequences of numbers (eg. Budescu, 1987; Rabinowitz, Dunlap, Grant, & Campione, 1989) and two-dimensional black and white grids (Falk, 1981). In producing binary sequences, people alternate with a probability of approximately 0.6, rather than the 0.5 that is seen in sequences produced by a random generating process. This preference for alternation results in subjectively random sequences containing less runs – such as an interrupted series of heads in a set of coin flips – than might be expected by chance (Lopes, 1982).

Theories of subjective randomness

A number of theories have been proposed to account for the accuracy of Reichenbach's conjecture. These theories have included postulating that people develop a concept of randomness that differs from the true definition of the term (eg. Budescu, 1987; Falk, 1981; Skinner, 1942), and that limited short-term memory might contribute to people's responses (Baddeley, 1966; Kareev, 1992; 1995; Wiegersma, 1982). Most recently, Falk and Konold (1997) suggested that the concept of randomness can be connected to the subjective complexity of a sequence, characterized by the difficulty of specifying a rule by which a sequence can be generated. This idea is related to a notion of complexity based on description length (Li & Vitanyi, 1997), and has been considered elsewhere in psychology (Chater, 1996).

The account of randomness that has had the strongest influence upon the wider literature of cognitive psychology is Kahneman and Tversky's (1972) suggestion that people may be attempting to produce sequences that are "representative" of the output of a random generating process. For sequences, this means that the number of elements of each type appearing in the sequence should correspond to the overall probability with which these elements occur. Random sequences should also maintain local representativeness, such that subsequences demonstrate the appropriate probabilities.

Formalizing representativeness

A major challenge for a theory of randomness based upon representativeness is to express exactly what it means for an outcome to be representative of a random generating process. One interpretation of this statement is that the outcome provides evidence for having been produced by a random generating process. This interpretation has the advantage of submitting easily to formalization in the language of probability theory.

If we are considering two candidate processes by which an outcome could be generated – one random, and one containing systematic regularities – the total evidence in favor of the random generating process can be assessed by the logarithm of the ratio of the probabilities of these processes

$$\log \frac{P(\text{random}|x)}{P(\text{regular}|x)}, \quad (1)$$

where $P(\text{random}|x)$ and $P(\text{regular}|x)$ are the probabilities of a random and a regular generating process respectively, given the outcome x .

This quantity can be computed using the odds form of Bayes' rule

$$\frac{P(\text{random}|x)}{P(\text{regular}|x)} = \frac{P(x|\text{random})}{P(x|\text{regular})} \frac{P(\text{random})}{P(\text{regular})}, \quad (2)$$

in which the term on the left-hand side of the equation is called the posterior odds, and the first and second terms on the right-hand side are called the likelihood ratio and prior odds, respectively. Of the latter two terms, the specific outcome x influences only the likelihood ratio. Thus the contribution of x to the evidence in favour of a random generating process can be measured by the logarithm of the likelihood ratio,

$$\text{random}(x) = \log \frac{P(x|\text{random})}{P(x|\text{regular})}. \quad (3)$$

This method of assessing the weight of evidence for a particular hypothesis provided by an observation is often used in Bayesian statistics, and the log likelihood-ratio given above is called a Bayes factor (Kass & Raftery, 1995). The Bayes factor for a set of independent observations will be the sum of their individual Bayes factors, and the expression has a clear information theoretic interpretation (Good, 1979). The above expression is also closely connected to the notion of minimum description length, connecting this approach to randomness with the ideas of Falk and Konold (1997) and Chater (1996).

Defining regularity

Evaluating the evidence that a particular outcome provides for a random generating process requires computing two probabilities: $P(x|\text{random})$ and $P(x|\text{regular})$. The first of these probabilities follows from the definition of the random generating process. For example, $P(\text{HHTHTTTH}|\text{random})$ is $(\frac{1}{2})^8$, as it would be for

any sequence of the same length. However, computing $P(x|\text{regular})$ requires specifying the probability of the observed outcome resulting from a generating process that involves regularities. While this probability is hard to define, it is in general easy to compute $P(x|h_i)$, where h_i might be some hypothesised regularity. In the case of sequences of heads and tails, for instance, h_i might correspond to a particular probability of observing heads, $P(\text{H}) = p$. In this case $P(\text{HHTHTTTH}|h_i)$ is $p^4(1-p)^4$. Using the calculus of probability, we can obtain $P(x|\text{regular})$ by summing over a set of hypothesized regularities, \mathcal{H} ,

$$P(x|\text{regular}) = \sum_{h_i \in \mathcal{H}} P(x|h_i)P(h_i|\text{regular}) \quad (4)$$

where $P(h_i|\text{regular})$ is a prior probability on h_i . In all applications discussed in this paper, we make the simplifying assumption that $P(h_i|\text{regular})$ is uniform over all $h_i \in \mathcal{H}$. However, we stress that this assumption is not necessary for the models we create, and the prior may in fact differ from uniformity in some realistic judgment contexts.

Random sequences

For the case of binary sequences, such as those that might be produced by flipping a coin, possible regularities can be divided into two classes. One class assumes that flips are independent, and the regularities it contains are assertions about the value of $P(\text{H})$. The second class includes regularities that make reference to properties of subsequences containing more than a single element, such as alternation, runs, and symmetries. Since this second class is less well defined, it is instructive to examine the account that can be obtained just by using the first class of regularities.

Taking \mathcal{H} to be all values of $p = P(\text{H}) \in [0, 1]$, we have $P(H, T|\text{random}) = (\frac{1}{2})^{H+T}$ and $P(H, T|\text{regular}) = \int_0^1 p^H (1-p)^T dp$, where H, T are the sufficient statistics of a particular sequence containing H heads and T tails. Completing the integral, it follows from (3) that

$$\text{random}(H, T) = \log \binom{H+T}{H} + f(H+T), \quad (5)$$

where $f(H+T)$ is $-\log 2^{H+T} - \log(H+T+1)$, a fixed function of the total number of flips in the sequence. This result has a number of appealing properties. Firstly, it is maximized when $H = T$, which is consistent with Kahneman and Tversky's (1972) original description of the representativeness of random sequences. Secondly, the ratio involved essentially measures the size of the set of sequences sharing the same number of heads and tails. A sequence like HHHHHHHH is unique in its composition, whereas HHTHTTTH has a composition much more commonly obtained by flipping a coin eight times.

The Zenith radio data

Having defined a framework for analyzing the subjective randomness of sequences, we have the opportunity to develop a specific model. One classic data set concerning

the production of random sequences is the Zenith radio data. These data were obtained as a result of an attempt by the Zenith corporation to test the hypothesis that people are sensitive to psychic transmissions. On several occasions in 1937, a radio program took place during which a group of psychics would transmit a randomly generated binary sequence to the receptive minds of their listeners. The listeners were asked to write down the sequence that they received, one element at a time. The binary choices included heads and tails, light and dark, black and white, and several symbols commonly used in tests of psychic abilities, and all sequences contained a total of five symbols. Listeners then mailed in their responses, which were analyzed. These responses demonstrated strong preferences for particular sequences, but there was no systematic effect of the actual sequence that was transmitted (Goodfellow, 1938). The data are thus a rich source of information about response preferences for random sequences. The relative frequencies of the different sequences, collapsed over choice of first symbol, are shown in the upper panel of Figure 1.

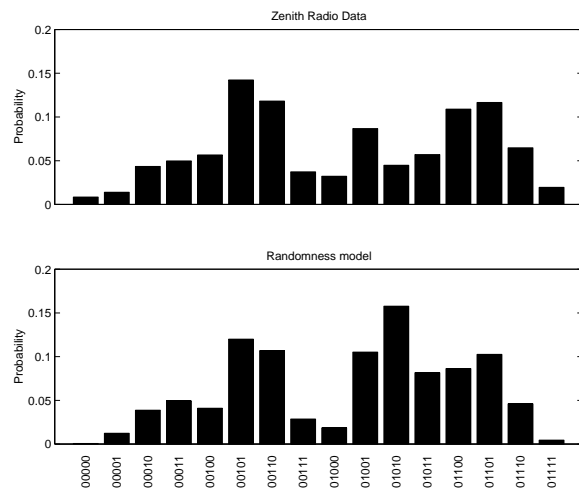


Figure 1: The upper panel shows the original Zenith radio data, representing the responses of 20,099 participants, from Goodfellow (1938). The lower panel shows the predictions of the randomness model. Sequences are collapsed over the initial choice, represented by 0.

Modeling random sequence production

One of the most important characteristics of the Zenith radio data is that people’s responses were produced sequentially. In producing each element of the sequence, people had knowledge of the previous elements. Kahneman and Tversky (1972) suggested that in producing such sequences, people pay attention to the local representativeness of their choices – the representativeness of each subsequence.

To capture this idea, we define L_k to be the local representativeness of choosing H as the k th response – the

extent to which H results in a more random outcome than T, assessed over the subsequences starting one step back, two steps back, and so forth,

$$L_k = \sum_{i=1}^{k-1} \text{random}(H_i + 1, T_i) - \text{random}(H_i, T_i + 1), \quad (6)$$

where the H_i, T_i are the tallies of heads and tails counting back i steps in the sequence. We can then convert this quantity into a probability using a logistic function, to give a probability distribution for the k th response, R_k :

$$\begin{aligned} P(R_k = H) &= \frac{1}{1 + e^{-\lambda L_k}} \\ &= \frac{1}{1 + \prod_{i=1}^{k-1} \left(\frac{T_i + 1}{H_i + 1} \right)^{-\lambda}} \\ &= \frac{\prod_{i=1}^{k-1} (T_i + 1)^\lambda}{\prod_{i=1}^{k-1} (T_i + 1)^\lambda + \prod_{i=1}^{k-1} (H_i + 1)^\lambda}. \end{aligned} \quad (7)$$

The λ parameter scales the effect that L_k has on the resulting probability. The probability of the sequence as a whole is then the product of the probabilities of the R_k , and the result defines a probability distribution over the set of binary sequences of length k . This distribution is shown in the lower panel of Figure 1 for $k = 5$.

This simple model provides a remarkably good account of the response preferences people demonstrated in the Zenith radio experiment. There is one clear discrepancy: the model predicts that the sequence 01010, equivalent to HTHTH or THTHT, should occur far more often than in the data. We can explain people’s avoidance of this sequence by the fact that alternation itself forms a regularity, which could easily be introduced into the hypothesis space. More striking is the account the model gives of the different frequencies of sequences with less apparent regularities, such as 00001 and 00010. Excluding the discrepant data point, the model gives a parameter-free ordinal correlation $r_s = 0.97$, and with $\lambda = 0.6$ has a linear correlation $r = 0.95$. Interestingly, the model predicts alternation, for sequences that are otherwise equally representative, with a probability of $\frac{1}{1+2^{-\lambda}}$. With the value of λ used in fitting the Zenith radio data, the resulting predicted probability of alternation is 0.6, consistent with previous findings (eg. Falk, 1981).

Pick a number

Research on subjective randomness has focused almost exclusively on sequences, but sequences are not the only stimuli that excite our intuitions about chance. In particular, random numbers loom larger in life than in the literature, although there have been a few studies that have investigated response preferences for numbers between 0 and 9. Kubovy and Pstotka (1976) reported the frequency with which people produce numbers between 0 and 9 when asked to pick a number, aggregated across several studies. These results are shown in the upper panel of Figure 2. People showed a clear preference for

the number 7, which Kubovy and Psotka (1976, p. 294) explained with reference to the properties of the numbers involved – for example, 6 is even, and a multiple of 3, but it is harder to find properties of 7. This explanation is suggestive of the kinds of regular generating processes that could be involved in producing numbers. Shepard and Arable (1979) found that similarity judgments about numbers could be captured by properties like those described by Kubovy and Psotka (1976), such as being even numbers, powers of 2, or occupying special positions such as endpoints.

Taking the arithmetic properties of numbers to constitute hypothetical regularities, we can specify the quantities necessary to compute $\text{random}(x)$. Our h_i are sets of numbers that share some property, such as the set of even numbers between 0 and 9. For any h_i , we define $P(x|h_i) = \frac{1}{|h_i|}$ for $x \in h_i$ and 0 otherwise, where $|h_i|$ is the size of the set. This means that observations generated from a regularity are uniformly sampled from that regularity. Setting $P(h_i|\text{regular})$ to give equal weight to all h_i , we can compute $P(x|\text{regular})$.

This model can be applied to the data of Kubovy and Psotka (1976). Since there are ten possible responses, we have $P(x|\text{random}) = \frac{1}{10}$. Taking hypothetical regularities of multiples of 2 ($\{0, 2, 4, 6, 8\}$), multiples of 3 ($\{3, 6, 9\}$), multiples of 5 ($\{0, 5\}$), powers of 2 ($\{2, 4, 8\}$), and endpoints ($\{0\}, \{1\}, \{9\}$), we obtain the values of $\text{random}(x)$ shown in the lower panel of Figure 2. Randomness also needs to be included in \mathcal{H} so that $\text{random}(x)$ is defined when x is not in any other regularity. Its inclusion is analogous to the incorporation of a noise process, and is in fact formally identical in this case. The order of the model predictions is a parameter free result, and gives the ordinal correlation $r_s = 0.99$. Applying a single parameter power transformation to the predictions, $y' = (y - \min(y))^{0.98}$, gives $r = 0.95$.

Coincidences

The surprising frequency with which unlikely events tend to occur has drawn attention from a number of psychologists and statisticians. Diaconis and Mosteller (1989), in their analysis of such phenomena, define a coincidence as ‘...a surprising concurrence of events, perceived as meaningfully related, with no apparent causal connection’ (p. 853). They go on to suggest that the “surprising” frequency of these events is due to the flexibility that we allow in identifying meaningful relationships. Together with the fact that everyday life provides a vast number of opportunities for coincidences to occur, our willingness to tolerate near misses and to consider each of a number of possible concurrences meaningful contributes to explaining the frequency with which coincidences occur. Diaconis and Mosteller suggested that the surprise that people show at the solution to the Birthday Problem – the fact that only 23 people are required to give a 50% chance of two people sharing the same birthday – suggests that similar neglect of combinatorial growth contributes to the underestimation of the




Figure 2: The upper panel shows number production data from Kubovy and Psotka (1976), taken from 1,770 participants choosing numbers between 0 and 9. The lower panel shows the transformed predictions of the randomness model.

likelihood of coincidences. Psychological research addressing coincidences seems consistent with this view, suggesting that selective memory (Hintzman, Asher, & Stern, 1978) and preferential weighting of first-hand experiences (Falk & MacGregor, 1983) might facilitate the under-estimation of the probability of events.

Not just likelihood...

The above analyses reflect the same bias that made it difficult to construct a probabilistic account of randomness: the notion that people’s judgments reflect the likelihood of particular outcomes. Subjectively, coincidences are events that seem unlikely, and are hence surprising when they occur. However, just as with random sequences, sets of events that are equally likely to be produced by a random generating process differ in the degree to which they seem to be coincidences. Following Diaconis and Mosteller’s suggestion that the Birthday Problem provides a domain for the investigation of coincidences, consider the kinds of coincidences formed by sets of birthdays. If we meet four people and find out that their birthdays are October 4, October 4, October 4, and October 4, this is a much bigger coincidence than if the same people have birthdays May 14, July 8, August 21, and December 25, despite the fact that these sets of birthdays are equally likely to be observed by chance. The way that these sets of birthdays differ is that one of them contains an obvious regularity: all four birthdays occur on the same day.

Modeling coincidences

Just as sequences differ in the amount of evidence they provide for having been produced by a random generating process, sets of birthdays differ in how much evi-

dence they provide for having been produced by a process that contains regularities. We argue that the amount of evidence that an event provides for a regular generating process will correspond to how big a coincidence it seems, and that this can be computed in the same way as for randomness,

$$\text{coincidence}(x) = \log \frac{P(x|\text{regular})}{P(x|\text{random})}. \quad (9)$$

To apply this model we have to define the regularities \mathcal{H} . For birthdays, these regularities should correspond to relationships that can exist among dates. Our model of coincidences used a set of regularities that reflected proximity in date (from 1 to 30 days), belonging to the same calendar month, and having the same calendar date (eg. January 17, March 17, September 17, December 17). We also assumed that each year consists of 12 months of 30 days each. Thus, for a set of n birthdays, $X = \{x_1, \dots, x_n\}$, we have $P(X|\text{random}) = (\frac{1}{360})^n$. In defining $P(X|\text{regular})$, we want to respect the fact that regularities among birthdays are still striking even when they are embedded in noise – for instance, February 2, March 26, April 3, June 12, June 12, June 12, November 22 still provides strong evidence for a regularity in the generating process. To allow the model to tolerate noisy regularities, we can introduce a noise term α into $P(X|h_i)$. The probability calculus lets us integrate out unwanted parameters, so the introduction of a noise process need not result in adding a numerical free parameter to the model. In particular, $P(X|h_i) = \int_0^1 P(X|\alpha, h_i)P(\alpha|h_i)d\alpha$. Assuming that the dates we observe are independent, we have $P(X|\alpha, h_i) = \prod_{x_j \in X} P(x_j|\alpha, h_i)$, and, taking a uniform prior on α , $P(X|h_i)$ is simply $\int_0^1 \prod_{x_j \in X} P(x_j|\alpha, h_i)d\alpha$, where

$$P(x_j|\alpha, h_i) = \begin{cases} \frac{\alpha}{360} + (1 - \alpha) \frac{1}{|h_i|} & x_j \in h_i \\ \frac{\alpha}{360} & x_j \notin h_i \end{cases}. \quad (10)$$

This corresponds to dates being sampled uniformly from the entire year with probability α , and uniformly from the regularity with probability $(1 - \alpha)$. The resulting $P(X|h_i)$ can then be substituted into (4), and taking a uniform distribution for $P(h_i|\text{regular})$ gives $P(X|\text{regular})$.

How big a coincidence?

The model outlined above makes strong predictions about the degree to which different sets of birthdays should be judged to constitute coincidences. We conducted a simple experiment to examine these predictions. The participants were 93 undergraduates from Stanford University, participating for partial course credit. Fourteen potential relationships between birthdays were examined, using two sets of dates. Each participant saw one set of dates, in a random order. The dates reflected: 2, 4, 6, and 8 apparently unrelated birthdays, 2 birthdays on the same day, 2 birthdays in 2 days across a month boundary, 4 birthdays on the same day, 4 birthdays in

one week across a month boundary, 4 birthdays in the same calendar month, 4 birthdays with the same calendar dates, and 2 same day, 4 same day, and 4 same date with an additional 4 unrelated birthdays, as well as 4 same week with an additional 2 unrelated birthdays. These dates were delivered in a questionnaire. Each participant was instructed to rate how big a coincidence each set of dates was, using a scale in which 1 denoted no coincidence and 10 denoted a very big coincidence.

The results of the experiment and the model predictions are shown in the top and middle panels of Figure 3 respectively. Again, the ordinal predictions of the model are parameter free, with $r_s = 0.94$. Applying the transformation $y' = (y - \min(y))^{0.48}$, gives $r = 0.95$. The main discrepancies between the model and the data are the four birthdays that occur in the same calendar month, and the ordering of the random dates. The former could be addressed by increasing the prior probability given to the regularity of being in the same calendar month – clearly this was given greater weight by the participants than by the model. Explaining the increase in the judged coincidence with larger sets of unrelated dates is more difficult, but may be a result of opportunistic coincidences: as more dates are provided, participants have more opportunities to identify complex regularities or find dates of personal relevance. This process can be incorporated into the model, at the cost of greater complexity.

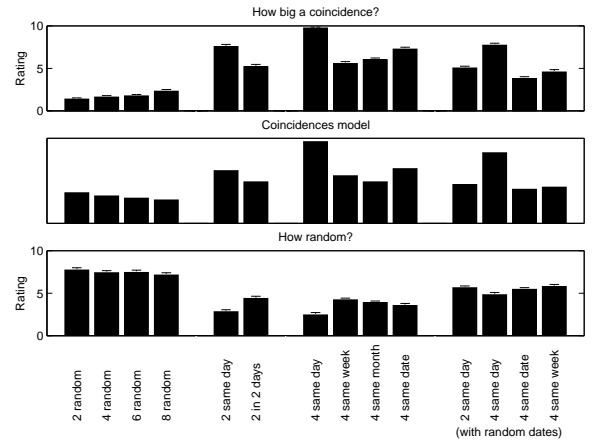


Figure 3: The top panel shows the judged extent of coincidence for each set of dates. The middle panel is the predictions of the coincidences model, subjected to a transformation described in the text. The bottom panel shows randomness judgments for the same stimuli.

Relating randomness and coincidences

Judgments of randomness and coincidences both reflect the evidence that a set of observations provides for having been produced by a particular generating process. Events that provide good evidence for a random generating process are viewed as random, while events that provide evidence for a generating process incorporating some regularity seem like coincidences. By examining

(3) and (9), we see that these phenomena are formally identified as inversely related: coincidences are events that deviate from our notions of randomness.

We conducted a further experiment to see if this relationship was borne out in people's judgments. Participants were 120 undergraduates from Stanford University, participating for partial course credit. The dates were the same as those used previously, and delivered in similar format. Each participant was instructed to rate how random each set of dates was, using a scale in which 1 denoted not at all random and 10 denoted very random.

The results of this experiment are shown in the bottom panel of Figure 3. The correlation between the randomness judgments and the coincidence judgments is $r = -0.94$, consistent with the hypothesis that randomness and coincidences are inversely linearly related. The main discrepancy between the two data sets is that the addition of unrelated dates seems to affect randomness judgments more than coincidence judgments.

Conclusion

The models we have discussed in this paper provide a connection between people's intuitions about chance, expressed in judgments about randomness and coincidences, and the formal structure of probability theory. This connection depends upon changing the way we model questions about probability. Rather than considering the likelihood of events being produced by a particular generating process, our models address the question of how much more likely a set of events makes a particular generating process. This is a structural inference, drawing conclusions about the world from observed data. Framed in this way, people's judgments are revealed to accurately approximate the statistical evidence that observations provide for having been produced by a particular generating process. The apparent inaccuracy of our intuitions may thus be a result of considering normative theories based upon the likelihood of events rather than the evidence they provide for a structural inference.

References

- Baddeley, A. D. (1966). The capacity of generating information by randomization. *Quarterly Journal of Experimental Psychology*, 18:119–129.
- Bar-Hillel, M. and Wagenaar, W. A. (1991). The perception of randomness. *Advances in applied mathematics*, 12:428–454.
- Budescu, D. V. (1987). A Markov model for generation of random binary sequences. *Journal of Experimental Psychology: Human perception and performance*, 12:25–39.
- Chaitin, G. J. (2001). *Exploring randomness*. Springer Verlag, London.
- Chater, N. (1996). Reconciling simplicity and likelihood principles in perceptual organization. *Psychological Review*, 103:566–581.
- Diaconis, P. and Mosteller, F. (1989). Methods for studying coincidences. *Journal of the American Statistical Association*, 84:853–861.
- Falk, R. (1981). The perception of randomness. In *Proceedings of the fifth international conference for the psychology of mathematics education*, volume 1, pages 222–229, Grenoble, France. Laboratoire IMAG.
- Falk, R. and Konold, C. (1997). Making sense of randomness: Implicit encoding as a bias for judgment. *Psychological Review*, 104:301–318.
- Falk, R. and MacGregor, D. (1983). The surprisingness of coincidences. In Humphreys, P., Svenson, O., and Vari, A., editors, *Analysing and aiding decision processes*, pages 489–502. North-Holland, New York.
- Good, I. J. (1979). A. M. Turing's statistical work in World War II. *Biometrika*, 66:393–396.
- Goodfellow, L. D. (1938). A psychological interpretation of the results of the Zenith radio experiments in telepathy. *Journal of Experimental Psychology*, 23:601–632.
- Hintzman, D. L., Asher, S. J., and Stern, L. D. (1978). Incidental retrieval and memory for coincidences. In Gruneberg, M. M., Morris, P. E., and Sykes, R. N., editors, *Practical aspects of memory*, pages 61–68. Academic Press, New York.
- Kac, M. (1983). What is random? *American Scientist*, 71:405–406.
- Kahneman, D. and Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3:430–454.
- Kareev, Y. (1992). Not that bad after all: Generation of random sequences. *Journal of Experimental Psychology: Human Perception and Performance*, 18:1189–1194.
- Kareev, Y. (1995). Through a narrow window: Working memory capacity and the detection of covariation. *Cognition*, 56:263–269.
- Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90:773–795.
- Kubovy, M. and Psotka, J. (1976). The predominance of seven and the apparent spontaneity of numerical choices. *Journal of Experimental Psychology: Human Perception and Performance*, 2:291–294.
- Li, M. and Vitanyi, P. (1997). *An introduction to Kolmogorov complexity and its applications*. Springer Verlag, London.
- Lopes, L. (1982). Doing the impossible: A note on induction and the experience of randomness. *Journal of Experimental Psychology*, 8:626–636.
- Rabinowitz, F. M., Dunlap, W. P., Grant, M. J., and Campione, J. C. (1989). The rules used by children and adults to generate random numbers. *Journal of Mathematical Psychology*, 33:227–287.
- Reichenbach, H. (1934/1949). *The theory of probability*. University of California Press, Berkeley.
- Shepard, R. and Arabie, P. (1979). Additive clustering: Representation of similarities as combinations of discrete overlapping properties. *Psychological Review*, 86:87–123.
- Skinner, B. F. (1942). The processes involved in the repeated guessing of alternatives. *Journal of Experimental Psychology*, 39:322–326.
- Tune, G. S. (1964). Response preferences: A review of some relevant literature. *Psychological Bulletin*, 61:286–302.
- Wagenaar, W. A. (1972). Generation of random sequences by human subjects: A critical survey of literature. *Psychological Bulletin*, 77:65–72.
- Wiegiersma, S. (1982). Can repetition avoidance in randomization be explained by randomness concepts? *Psychological Research*, 44:189–198.

Acknowledgments

This work was supported by a Hackett studentship to the first author and funds from Mitsubishi Electric Research Laboratories. The authors thank Persi Diaconis for inspiration and conversation, Tania Lombrozo, Craig McKenzie and an anonymous reviewer for helpful comments, and Kevin Lortie for finding the leak.