

# Stochastic Independence and Causal Connection

Michael Strevens

Draft of March 2011

## Abstract

Assumptions of stochastic independence are crucial to statistical models in science. But under what circumstances is it reasonable to suppose that two events are independent? When they are not causally or logically connected, so the usual story goes. But scientific models frequently treat causally dependent events as stochastically independent, raising the question whether there are kinds of causal connection that do not undermine stochastic independence. This paper provides one piece of an answer to this question, treating the simple case of two tossed coins with and without a midair collision.

## 1. The Importance of Stochastic Independence

Without pervasive stochastic independence, statistical reasoning would be of little use to us. The principle of total evidence tells us to take everything relevant into account when reasoning inductively, as a consequence of which, if we are thinking about an event  $e$ , the probability of  $e$  that matters for us is  $P(e|k)$ , where  $k$  is all our background knowledge. Were this probability to have different values for every possible  $k$ —for every possible set of background knowledge—we would have to compile an immense compendium of statistical knowledge in order to comply with the total evidence principle. This is true whether the probability in question is subjective, epistemic, or physical.

Fortunately, we seem to have good reason to regard most of our background knowledge as probabilistically irrelevant to many of the events about which we wish to reason statistically, or equivalently, many such events are stochastically independent of most of the background. In such cases, almost all of the background can be ignored; the probability distribution over  $e$  and its sibling events is consequently within our human cognitive grasp (Harman 1973; Pollock manuscript).

Statistical models in science illustrate this observation by imposing physical probability distributions over phenomena that take into account only a few parameters, and thus that assume implicitly that the events they describe are stochastically independent of all properties not captured by these parameters.<sup>1</sup> Kinetic theory's velocity distribution over gas molecules, for example, supposes that a molecule's velocity is independent of everything except temperature.<sup>2</sup> Statistical representations of certain kinds of experimental error assume that errors are uncorrelated, that is, that one error is independent of any other, as do "error functions" representing "noise" in many statistical models in the social sciences. Finally, stochastic models in population genetics make a number of independence assumptions, most notably, that one instance of a gene's being replicated in the next generation is independent of any other.

When is it possible to assume that two events or states of affairs are stochastically independent? And why does independence exist, when it does? To answer these questions, asked of the physical probability distributions ascribed by scientific theories, is of great importance for the foundations of physical probability. The matter has, however, received relatively little

---

1. Science's physical probability distributions are in the first instance over event types rather than event tokens; in what follows, I use the term "event" to mean event type, although much of what I say remains true when said of singular events.

2. A more careful statement of the independence assumption would exclude quantities logically related to velocity and temperature from the independence claim: kinetic theory does not, for example, assume that a molecule's velocity is independent of its momentum. But there is no need to get held up over such matters here.

attention in the literature.

In a notable exception to this neglect, Pollock (manuscript) derives the following principle from his theory of “nomic probability”: it is always defeasibly reasonable “to expect, in the absence of contrary information, that [events] will be statistically independent of one another” (p. 11).<sup>3</sup> You might of course find this thesis quite acceptable even if you do not share Pollock’s interpretation of probability, and many, perhaps the vast majority, of scientists seem to regard it as, at the very least, a valid heuristic. The usefulness of Pollock’s principle depends, however, on how common it is to have contrary information, and ultimately, how much contrary information there is to be had: if contrary information exists for every pair of properties or events, then a well-informed scientist will be unable to use the principle at all. What facts, then, will defeat the inference endorsed by the principle?

Accepting the received wisdom, Pollock remarks that either a causal connection or a logical connection between events will undermine the assumption of independence. You might, then, formulate a more informative heuristic as follows:

Expect, in the absence of contrary information, that events that are not causally or logically connected will be statistically independent of one another.

The “in the absence of contrary information” rider is still necessary because there are defeaters for independence not made explicit in the principle, the simplest example being statistical evidence against independence. Nevertheless, the revised principle is sufficiently explicit for my purposes.

Here is the problem that will motivate the rest of the paper: often we want to assume the independence of causally connected events. Most often, the connection is by way of a common causal origin: the events have among

---

3. Pollock says “properties” rather than “events”. I remind you that by “event” I mean “event type”; with this understanding, I think that Pollock would allow my reformulation.

their causes numerically the same event. The existence of a common cause frequently establishes a correlation; that barometer readings and the weather are both caused in part by properties of atmospheric pressure, for example, accounts for the statistical connection between barometer drops and storms. Thus, a shared etiology can undermine independence.

Events with a common cause are, nevertheless, frequently treated as independent. There are two kinds of cases worth mentioning. First, on a liberal conception of what it is to share a common cause, almost any two events with some spatiotemporal proximity have something causal in common. On a very liberal conception, everything is perhaps causally connected by the Big Bang.

Second, even if you dismiss such worries as overly sensitive to minor causal commonalities, there are situations where the commonality is anything but minor. Even a case as iconic as repeated coin tosses furnishes an example: that consecutive coin tosses are made by the same person in the same frame of mind surely constitutes a significant similarity in causal history. Why think that the outcome of one such toss provides no information about the outcome of the other? Their common origin surely cannot be discounted on the grounds of causal unimportance: the croupier who tosses the coin determines almost everything about the initial conditions that is relevant to the outcome of a toss.

Statistical models in science provide many more examples. The position and velocity of a gas molecule are treated as independent of the positions and velocities of the other molecules, even though they are determined entirely by interaction with those other molecules. Similarly, an instance of gene replication in population genetics is treated as independent of the other instances, even though replication is determined in all cases by events in a single, densely causally connected ecosystem.

These considerations together suggest that stochastic independence can and often does exist in spite of a substantial, even an entire, overlap in causal

origins. Further, in science and in everyday life we take great advantage of this fact. Pollock's principle and other such heuristics, valid though they might be, cannot explain why certain causal connections do not undermine independence nor when we are justified in assuming independence in spite of causal commonalities.

When the physical probability distributions in question are irreducible, as the distributions of quantum mechanics are often supposed to be, the answers may be spelled out in the fundamental laws of nature. But the great majority of physical probabilities in science and in everyday life—the probabilities attached to gambling setups, actuarial probabilities, the probabilities of statistical physics and evolutionary biology, the probabilities discussed by Harman and Pollock—are not irreducible. This paper will focus on independence in these reducible distributions, that is, in physical probability distributions that are presumably grounded in, among other things, facts about the causal generation of the outcomes.<sup>4</sup>

Here is a problem, then, in urgent need of philosophical attention: state the conditions under which causally connected events may be assumed to be independent, relative to the reducible physical probability distributions imposed on them by our best scientific models and theories. Or, to take a slightly different approach, in conditions that otherwise favor independence, state the kinds of causal interaction that do and do not undermine independence.

The extent of the project is vast. In this paper, I put just a small piece in place, deriving conditions for independence among the outcomes of coin tosses causally connected in two ways: first, two consecutive coin tosses made by a single croupier, and second, two simultaneous coin tosses that collide in

---

4. Some philosophers deny the possibility of reducible physical probability. What they cannot deny is the existence of powerful scientific models that represent deterministically produced phenomena using probability distributions that incorporate sweeping independence judgments. The success of this practice must be explained even if reducible probability is in some metaphysical sense not "real" probability.

midair. My results embrace and extend previous work in Strevens (2003) (for an overview of which, see Strevens (2005)). The case of colliding and other tossed coins is of course of limited interest in itself, but some more general lessons about causation and independence will be clear enough.

## 2. The Probabilistic Dynamics of Coin Tosses

Consider the following simple coin tossing setup: a coin is tossed with a variable rotational velocity, allowed to spin for a fixed time, and then stopped. Whichever face is uppermost determines the outcome of the toss, head or tails. Assuming that the underlying dynamics are deterministic, the outcome of such a toss is fully determined by the physics of the coin and a single variable initial condition, the coin's spin speed. (The physics of a real coin toss is rather more complex than this (Keller 1986; Diaconis et al. 2007), but the additional complications that would arise from a more realistic model add nothing to the understanding of independence.)

Consider a function that maps any possible initial spin speed for such a toss to the outcome produced. The spin speeds are real numbers and the outcome will be represented by an integer, zero for tails and one for heads. ("Edge" is ignored.) I call this function the coin's *evolution function*; a typical shape for the evolution function is show in in figure 1. To help you see its shape, the area under the function is shaded.




Figure 1: The evolution function for obtaining heads on a simple coin toss with initial spin speed  $v$

The simple coin toss's evolution function has two properties that it shares with the evolution functions for many other classic gambling setups:

1. The outcome varies from tails to heads and back again very quickly as the value of the spin speed  $v$  changes, or in other words, very small changes in the initial spin speed of a coin will change the outcome of a toss from heads to tails and vice versa. (This sort of sensitivity to initial conditions will, you might think, amplify the tendency of a common causal history to undermine independence.)
2. For any small but not-too-small interval of  $v$ , the ratio of values of  $v$  producing heads to those producing tails is the same (in the case of the coin, one-to-one, that is, one-half). Or better:  $v$  can be partitioned into small intervals, in each of which the ratio of heads- to tails-producing values of  $v$ —the ratio of gray to white sections of the graph—is the same.

I call the conjunction of these two properties *microconstancy*, and I call the constant ratio of heads to tails the *strike ratio* for heads. (Formal definitions are given in Strevens (2003).)<sup>5</sup>

Many philosophers and scientists writing about the foundations of physical probability—von Kries, Poincaré, Hopf, and others—have thought that the microconstancy of the coin's evolution function with strike ratio one-half explains the fact that the physical probability of heads is one-half.<sup>6</sup> I fully agree (Strevens forthcoming).

---

5. The microconstancy of an evolution function is always relative to an outcome. It is also relative to a measure, a way of quantifying a setup's initial conditions. The correct measure to use in the present framework is the measure with respect to which the probability distribution over the setup's initial conditions is stated. In what follows I take the initial condition distribution, and thus the measure, as given.

6. For a guide to the older work, see von Plato (1983), and for more recent work, Diaconis et al. (2007). Reichenbach (2008) and Hopf (1934, §9) briefly explore the relationship between microconstancy and independence, deriving a result much like that presented in the next section. For the differences between the approaches taken by the earlier writers and my own approach, see Strevens (2003, §2.A) and Strevens (forthcoming).

For my purposes in this paper, however, it is not necessary to take a position on the metaphysics of the probability of heads. Rather, I make the following assumptions. First, in any coin toss, there is a physical probability distribution over the initial spin speed  $v$ . Second, the probability of heads is equal to the probability that the initial speed distribution assigns to heads-producing values of  $v$ .<sup>7</sup>

From these posits, it is possible to gain considerable insight into a problem posed above: why are the outcomes of consecutive coin tosses independent, when they have a common causal origin, namely, a single human croupier?

The key to the insight is a simple mathematical theorem which also spurred the effort to ground physical probabilities in the property of microconstancy: if the evolution function of some setup is microconstant with strike ratio  $p$  for an outcome  $e$ , then any initial condition distribution, provided that it has a certain smoothness property, will induce the same probability for  $e$ , equal to  $p$ .


The smoothness in question can be variously defined. One way to think about it is as follows. Microconstancy consists in the initial condition space's being divisible into small contiguous (i.e., connected) regions each with the same proportion of  $e$ -producing initial conditions. An initial condition distribution will induce a physical probability equal to this proportion if it is approximately uniform across all of these small regions. Call this kind of smoothness—sufficient but not necessary for the equality of strike ratio and probability—*macroperiodicity*.

The theorem—that any macroperiodic probability distribution over the initial conditions of a microconstant setup induces a physical probability for an outcome equal to its strike ratio—is easily appreciated in pictures. Figure 2 shows two distributions over a coin's spin speed  $v$ ; in each case,

---

7. In Strevens (forthcoming) I endorse the second of these assumptions (in a back-handed way; see the end of §4), but also show how to ground a probability for heads without the first assumption. In the present paper, then, the first assumption is made for purely expository reasons.

the area under the distribution that is shaded gray, corresponding to the probability ascribed by the distribution to heads-producing values of  $\nu$ , is equal to the strike ratio for heads of one-half. The result generalizes to setups



*Figure 2:* Different macroperiodic distributions over initial spin speed  $\nu$  induce the same probability for heads, equal to the strike ratio for heads

with more than one initial condition variable (and in fact can be somewhat strengthened; Strevens (2003, §2.25) states the stronger result and provides formal proofs of all results stated here).

You can see from the theorem (or the figure) why two croupiers, physically different enough that they impose a different physical probability distributions over the spin speeds that they impart to their coins, may nevertheless induce equal probabilities for heads.<sup>8</sup> Add a big empirical assumption to the mix—that probability distributions over spin speed tend to be macroperiodic—and you have an explanation why most tossed coins land heads about one-half of the time. These explanations are discussed further in Strevens (2003) and Strevens (forthcoming), and also in the literature indirectly referenced in note 6. Rather than go into the details any further

---

8. Or at least, probabilities that are almost equal.

here, however, let me turn to the question of the independence of consecutive tosses.

### 3. Consecutive Tosses by the Same Croupier

Consider two consecutive tosses by the same croupier. Why, despite their common causal origin, are the outcomes of the tosses independent?

You might expect a failure of independence for the following reason (foreshadowed above). Different croupiers, I will suppose, toss the coin differently enough that they impose different physical probability distributions over initial spin speed. You might therefore suppose that two tosses from the same croupier will fail to be independent because information about the outcome of one will give you information about that particular croupier's initial spin speed distribution. For reasons given in the last section, however, this is not so: if all croupiers' distributions are macroperiodic, then learning the outcomes of tosses tells you nothing that discriminates one croupier's distribution from another.

This defuses a reason for thinking that independence is undermined, but a far stronger conclusion is possible: it can be shown that, given the right assumptions about macroperiodicity, any two tosses, including any two consecutive tosses from the same croupier, will be independent for sure.

Why? Performing two consecutive tosses can be thought of as a single "conjoint" process that takes as its initial conditions the spin speeds of the two tosses and produces as its outcome an ordered pair containing the outcomes of the two tosses. If the evolution function for an individual toss is microconstant with a strike ratio of one-half for heads, then the evolution function for (say) a pair of heads on the conjoint setup is microconstant with a strike ratio equal to the product of the strike ratios for the individual outcomes, namely, one-quarter. Such an evolution function is shown in figure 3; comparison of this graph with figure 1 should make it clear why the result holds.

It follows that, if the joint density over the initial conditions of the two

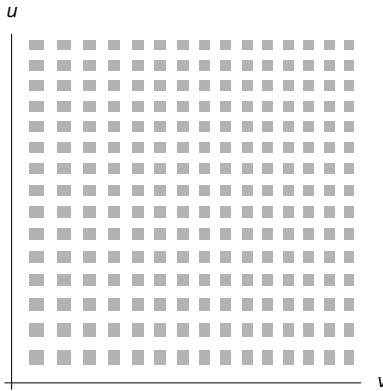


Figure 3: Conjoint evolution function for the event of obtaining heads on two consecutive coin tosses with spin speeds of  $u$  and  $v$ ; the shaded regions show values of  $u$  and  $v$  that produce two heads

tosses is macroperiodic, then the probability for any conjoint outcome is equal to the product of the probabilities for the single outcomes, as required for independence. (The joint density represents the probability that two consecutive tosses have any two particular spin speeds.)

This result is entirely general. Take two microconstant setups, one with a strike ratio of  $p$  for outcome  $e$  and the other with a strike ratio of  $q$  for event  $f$ , and the conjoint evolution function for pairs of trials on the setups will itself be microconstant with a strike ratio for the conjoint outcome  $(e, f)$  of  $pq$ . If the joint density is macroperiodic, then the probability for the conjoint outcome is equal to the product of the probabilities for the single outcomes:  $P(e f) = P(e)P(f)$ .<sup>9</sup> That is necessary and sufficient for independence.

The condition required of the initial conditions by this independence result—macroperiodicity of the joint distribution—is much weaker than independence.<sup>10</sup> Thus, you can have a certain degree of correlation between

9. Here I also assume, of course, that the densities for individual trials are macroperiodic, so that the probabilities are equal to the trials' strike ratios.

10. Mathematically speaking, almost all macroperiodic joint densities represent some degree of correlation between the events whose probabilities they encode.

the initial speeds of two consecutive coin tosses, representing perhaps the croupier's state of mind or other ephemeral physiological or environmental conditions, without losing independence in the outcomes. The microconstant dynamics in effect breaks down the correlation; it is an independence-creating machine, nullifying the correlating power of a common causal origin.

This observation goes a long way, I think, towards explaining why independence is widespread even among events that share much of their causal history. (More would have to be said, of course, to establish the existence of microconstancy in a wide range of cases of interest.)

#### 4. Colliding Coins

The independence result for consecutive tosses stated in the previous section assumes that, although the two trials in question may have a common causal origin, they do not interact in any way once their initial conditions are determined. (Mathematically, this manifests itself as the assumption that the trials share no initial conditions.) Sometimes, however, independence is found even when such interaction is present. The probabilities of population genetics, for example, represent the events over the span of a generation; within the generation, there is in most cases considerable potential for interaction between the gene-bearers. How can there be independence in spite of the interaction?<sup>11</sup>

In this final section, I want to start on an answer to this question by examining a case in which there is a causal interaction between two tossed coins: they collide in midair. I show that in some circumstances of practical

---

11. In the treatment of tossed coins in this section, I begin with probabilities for causally isolated coin tosses—coins that do not collide—and show that most collisions do not undermine independence. That method cannot be applied directly to the probabilities of population genetics, which are inherently extrinsic, and so do not exist in a causally isolated form. The way to demonstrate the independence of the population genetic probabilities is to build them out of probabilities for smaller causal steps that do come in an isolated form, as explained in Strevens (2003, chap. 4).

interest, the collision does not interfere with independence.

To keep things simple, I assume that the only effect of the collision is a change in the coins' spin speed, thus that after the collision, they continue to spin around the same axis. There is no transition to a more complex tumbling motion, then, though the methods I use are amenable to such a generalization.

In what follows, I will use, without proving, a result stated and proved in Strevens (2003), §3.6 and §3.B6:

The independence of outcomes produced by simultaneous microconstant trials (with macroperiodic individual and joint initial condition distributions) is not undermined by a causal interaction, provided that (a) the interaction occurs at the beginning of the trial, and (b) the interaction can be represented by an instantaneous, non-deflationary, microlinear transformation of the initial conditions.

A graphical sketch of the proof is provided in Strevens (2005).

Some definitions: A transformation is *non-deflationary* if it transforms any region into a region of equal or greater size. A simple example is  $y = 2x$ , which doubles the size of any interval. A transformation is *microlinear* if it can be approximated by a mosaic of linear transformations, or in other words, if for any small, contiguous region of the transformation's domain, there exists a linear transformation that has an approximately equivalent effect on the region (mapping the same points to roughly the same points). Microlinearity is of course relative to some standard for what is "micro", that is, what counts as a small region. In the theorem, the standard is set by the relevant evolution functions' microconstancy: loosely speaking, a region is small if it is the same size, or smaller, than the regions of constant proportion in virtue of which the evolution functions are microconstant.<sup>12</sup>

---

12. More technically, the microlinearity of a transformation can be defined as relative to a partition of its domain into connected sets: it is microlinear if it is approximately linear

The above result guarantees that the independence of simultaneously tossed coins will not be undermined by a midair collision if that collision effects a transformation of the initial conditions that is non-deflationary, microlinear, and occurs at the very beginning of the tosses. The last of these conditions is by far the most unwelcome.<sup>13</sup> My aim in the rest of this section is to remove the restriction, showing that the result holds for interactions at any point in the trial.

The structure of the argument is as follows: I show that under certain conditions a microlinear, non-deflationary transformation of a coin's spin speed partway through a toss is equivalent to a different transformation on its initial spin speed that is also microlinear and non-deflationary. Since the latter sort of transformation preserves independence, so does the former.

Suppose, then, that two coins are tossed, with spin speeds of  $u$  and  $v$  respectively, in such a way that they both spin for time  $T$ . At some time  $t_1$  less than  $T$ , the coins collide, effecting a change in their spin speeds (but as noted above, no change in their axes of rotation). Assume that a coin's spin speed does not significantly decrease over the course of the toss. (This is a significant assumption; I will revisit it later.)

Let me focus on the dynamics of the first coin. Supposing that spin speeds are represented in revolutions per time unit, the number of revolutions made by the first coin at the time of the collision  $t_1$  is  $t_1u$ . If the speed of the coin is altered at  $t_1$  to  $u'$ , then the number of revolutions the coin makes between  $t_1$  and  $T$  is  $t_2u'$ , where  $t_2 = T - t_1$ . Thus the total number of revolutions  $n$  made by the coin, which determines the outcome of this toss, is  $t_2u' + t_1u$ .

I will first consider the case where the transformation effected by the interaction is linear, and so can be represented as follows:

$$u' = au + bv + c$$

---

over any member of the partition. What the theorem above requires is microlinearity relative to a partition into sets of equal strike ratio.

13. For technical reasons, it created no problems for Strevens (2003).

for some  $a$ ,  $b$ , and  $c$ .<sup>14</sup> The total number of revolutions made by the coin is the sum of the revolutions made before the collision and the revolutions made after the collision:

$$\begin{aligned} n &= t_2 u' + t_1 u \\ &= t_2 (au + bv + c) + t_1 u \\ &= (at_2 + t_1)u + (bt_2)v + ct_2. \end{aligned}$$

My aim is to find a linear or linear plus constant transformation of  $u$  and  $v$  that yields the same formula for  $n$  when applied to  $u$  at the very beginning of the toss. This new transformation, if it exists, has the form:

$$du + ev + f.$$

If a transformation of this form is applied at the beginning of a toss, the total number of revolutions made by the coin is:

$$\begin{aligned} n &= T(du + ev + f) \\ &= dTu + eTv + fT. \end{aligned}$$

Such a transformation is equivalent to the actual transformation if there are values of  $d$ ,  $e$ , and  $f$  for which these two formulas for  $n$  are equivalent. Clearly, there are:

$$\begin{aligned} d &= (at_2 + t_1)/T \\ e &= bt_2/T \\ f &= ct_2/T. \end{aligned}$$

So the effect of a linear interaction on  $u$  partway through a coin toss is equivalent to the effect of a linear interaction—though a different one—at

---

14. A non-zero value for  $c$  would be physically peculiar, but since it is easy to handle, I do not rule it out here.

the beginning of the toss. The same is true for  $v$ . Thus the transformation of  $u$  and  $v$  partway through the toss is equivalent to a linear transformation of  $u$  and  $v$  at the beginning of the toss.

Under what conditions will this equivalent interaction at the beginning of the toss be non-deflationary? The “Further Proofs” section at the end of this paper shows that, provided certain physically plausible conditions hold, if the transformation effected by the actual interaction of the coins is non-deflationary, then the equivalent beginning-of-the-toss transformation is also non-deflationary.

Now consider the case in which the transformation effected by the interaction is microlinear, that is, linear over any small, contiguous region of its domain. From the reasoning above, over any such region, the effect of the transformation is equivalent to a different linear interaction at the beginning of the trial applied to the same region.

It follows that the effect of a microlinear interaction partway through a coin toss, which is made up of different linear interactions over different small regions, is approximately equivalent to the effect of a corresponding set of linear interactions over the same regions at the beginning of a coin toss, thus is equivalent to a microlinear interaction at the beginning of the toss. Further, provided that the original interaction is non-deflationary, the new interaction is also non-deflationary. Such an interaction will therefore preserve independence.

The argument as presented so far makes the rather narrow (if in many cases realistic) assumption that a coin’s spin speed remains effectively constant from the beginning to the end of the toss. The assumption may, however, be relaxed, as I will now explain.

The principal part of the derivation above is the demonstration that a linear transformation of spin speeds partway through a toss is equivalent to a linear transformation at the beginning of the toss. Inspection of the proof shows that what matters for the derivation is that the formula for the number

of revolutions made by the coin over a specified period of time, given an initial speed of  $u$ , is a linear function of  $u$ . This is true provided that the spin speed of the coin at any time  $t$ , given an initial speed of  $u$ , can be written  $f(t)u$  (where  $f(\cdot)$  need not itself be linear), since the number of revolutions made over a period of length  $t$  is then  $F(t)u$ , where  $F(t) = \int_0^t f(x) dx$ . To generalize the proof above, then, simply substitute  $F(t_1)$  for  $t_1$ ,  $F(t_2)$  for  $t_2$ , and  $F(T)$  for  $T$ . The linearity condition is satisfied by a number of possible dynamics for spin speed, such as exponential decay (on which the speed at time  $t$ , given an initial speed  $u$ , is  $ue^{-kt}$ ).

Two other issues must be addressed to secure independence. First, in order to have microlinearity of the beginning-of-the-toss transformation, it is not enough that the actual interaction transformation be microlinear, that is, linear over any small region of the speed space. It must be linear over the *inverse image* of any such region with respect to the speed evolution function (the function that determines how speed changes with time). A sufficient condition for this, given the microlinearity of the actual interaction transformation, is that the inverse image of any region of the speed space be at least as large as that region, because then we are guaranteed linearity over these “at least as large” sets, hence microlinearity. This sufficient condition amounts to the requirement that the speed evolution function be non-inflationary, or in other words, that a coin’s spin speed not increase over time. This condition is of course eminently reasonable for a tossed coin.

The second condition for independence is that the beginning-of-the-toss transformation be non-deflationary. The “Further Proofs” show that the non-inflation requirement on speed evolution is sufficient for the satisfaction of this condition (given the other conditions already in place).

In short, then, the independence result can be extended to any case in which the speed evolution function is linear in the speed and non-inflationary. You can likely see the prospect for further generalization to speed evolution functions that are microlinear, but let me leave that to another time.

## 5. Conclusions

Thanks to microconstancy, there is ample independence to be found among the outcomes of coin tosses. First, microconstancy eliminates correlations due to common causal origins, as when the initial spins of consecutive tosses are correlated due to the passing moods of the croupier. Second, microconstancy protects independence against a range of causal interactions between the coins once tossed, in particular those that are non-deflationary and microlinear, provided that the speed of the coin does not decrease significantly over the course of the toss.<sup>15</sup>

The interest of the results for coins depends, of course, on the extent of microconstancy, and in particular, on the question whether microconstancy is to be found in the kinds of systems concerning which scientific statistical models make broad assumptions of independence.

I have argued for the prevalence of microconstancy elsewhere (Strevens 2003). Even so, the results referenced and developed in this paper at best constitute only the beginnings of an extensive and important philosophical project, of understanding the reasons for physical probabilistic independence among causally connected outcomes. Some further thoughts about causation and independence can be found in the material supplementary to my book *Bigger than Chaos* posted at [www.strevens.org/chaos/](http://www.strevens.org/chaos/) under “Additions”. Those programmatic remarks show just how much work remains to be done.

---

15. Two additional, physically plausible, assumptions are made in the course of the derivation (see the “Further Proofs”):  $a$  and  $g$  are positive and  $b$  and  $h$  have the same sign.

### Further Proofs: Sufficient Conditions for Non-Deflation

The following proofs justify the claim that if a midair collision between two tossed coins effects a non-deflationary linear transformation of the spin speeds, then the equivalent transformation at the beginning of the tosses will also be non-deflationary, provided that some additional, physically plausible conditions hold.

Some new notation will simplify the algebra. Define  $\hat{t}_1$  and  $\hat{t}_2$  as the proportions of the total spin time that elapse before and after the collision, so that

$$\hat{t}_1 = t_1/T \quad \text{and} \quad \hat{t}_2 = t_2/T.$$

Then the coefficients  $d$  and  $e$  of the beginning-of-the-toss transformation of  $u$  equivalent to the actual interaction can be written

$$\begin{aligned} d &= a\hat{t}_2 + \hat{t}_1 \\ e &= b\hat{t}_2. \end{aligned}$$

The same notation can be used for the coefficients of the beginning-of-the-toss transformation of  $v$ . If the actual linear transformation of  $v$  effected by the collision is:

$$v' = gv + hu + i.$$

then the beginning-of-the-toss transformation of  $v$  equivalent to the actual interaction is

$$(g\hat{t}_2 + \hat{t}_1)v + h\hat{t}_2u + i\hat{t}_2$$

The complete beginning-of-the-toss transformation—the transformation on the  $u \times v$  space—is non-deflationary just in case its determinant is greater than or equal to 1. Its determinant is

$$\det \begin{pmatrix} a\hat{t}_2 + \hat{t}_1 & b\hat{t}_2 \\ h\hat{t}_2 & g\hat{t}_2 + \hat{t}_1 \end{pmatrix} = ag(\hat{t}_2)^2 + a\hat{t}_1\hat{t}_2 + g\hat{t}_1\hat{t}_2 + (\hat{t}_1)^2 - bh(\hat{t}_2)^2$$

$$= (ag - bh)(\hat{t}_2)^2 + \frac{(a + g)}{2} 2\hat{t}_1\hat{t}_2 + (\hat{t}_1)^2.$$

Since  $(\hat{t}_2)^2 + 2\hat{t}_1\hat{t}_2 + (\hat{t}_1)^2 = (\hat{t}_1 + \hat{t}_2)^2 = 1$ , a sufficient condition for the determinant to be greater than or equal than 1 is that the coefficients of  $(\hat{t}_2)^2$  and  $2\hat{t}_1\hat{t}_2$  in the above expression be greater than or equal to 1, that is, that

1.  $ag - bh \geq 1$ , and
2.  $(a + g) \geq 2$ .

The quantity  $ag - bh$  is the determinant of the actual transformation, that is, the transformation actually effected by the midair collision partway through the toss; it is greater than or equal to 1 just in case the original transformation is non-deflationary.

If it is assumed that  $a$  and  $g$  are positive (i.e., a coin's pre-interaction spin speed makes a positive rather than a negative contribution to its post-interaction spin speed) and that (because of physical symmetry) the signs of  $b$  and  $h$  are the same, then (1) entails that  $ag \geq 1$ , which in turn entails (2). Under these assumptions, then, if the interaction transformation partway through the toss is non-deflationary, its equivalent at the beginning of the toss is also non-deflationary.

Next consider the case where the coin's spin speed changes significantly over time. As in the main text, I assume that the speed at any time is a linear function of the initial speed: there exists a function  $f(\cdot)$  (not necessarily linear) such that, after time  $t$ , the speed of a coin with initial speed  $u$  will be  $f(t)u$ . Redefine  $\hat{t}_1$  and  $\hat{t}_2$  so that

$$\hat{t}_1 = F(t_1)/F(T) \quad \text{and} \quad \hat{t}_2 = F(t_2)/F(T).$$

where  $F(t) = \int_0^t f(x) dx$ . Then the relevant beginning-of-the-toss transformation can then be written just as I have written it above.

It can no longer be assumed that  $\hat{t}_1 + \hat{t}_2 = 1$ . However, if  $f(\cdot)$  is non-inflationary (spin speed remains the same or decreases over time), then as

demonstrated below,  $\hat{t}_1 + \hat{t}_2 \geq 1$ , which is sufficient for the above reasoning to apply to the more general case, so that the beginning-of-the-toss transformation is non-deflationary if the actual transformation is non-deflationary and the other conditions stated above apply.

To show that  $\hat{t}_1 + \hat{t}_2 \geq 1$ , first observe that if spin speed evolution is non-inflationary (i.e.,  $f(\cdot)$  is a non-increasing function) then

$$\int_0^a f(t) dt \geq \int_b^{b+a} f(t) dt$$

It then follows from the definition of  $F(\cdot)$  that

$$\begin{aligned} \hat{t}_1 + \hat{t}_2 &= \frac{F(t_1) + F(t_2)}{F(T)} \\ &= \frac{\int_0^{t_1} f(t) dt + \int_0^{t_2} f(t) dt}{\int_0^T f(t) dt} \\ &\geq \frac{\int_0^{t_1} f(t) dt + \int_{t_2}^T f(t) dt}{\int_0^T f(t) dt} \\ &\geq 1. \end{aligned}$$

## References

- Diaconis, P., S. Holmes, and R. Montgomery. (2007). Dynamical bias in the coin toss. *SIAM Review* 49:211–235.
- Harman, G. (1973). *Thought*. Princeton University Press, Princeton, NJ.
- Hopf, E. (1934). On causality, statistics and probability. *Journal of Mathematics and Physics* 13:51–102.
- Keller, J. (1986). The probability of heads. *American Mathematical Monthly* 93:191–197.
- von Plato, J. (1983). The method of arbitrary functions. *British Journal for the Philosophy of Science* 34:37–47.
- Pollock, J. (Manuscript). Probable probabilities.
- Reichenbach, H. (2008). *The Concept of Probability in the Mathematical Representation of Reality*. Translated and edited by F. Eberhardt and C. Glymour. Open Court, Chicago. Reichenbach's doctoral dissertation, originally published in 1916.
- Strevens, M. (2003). *Bigger than Chaos: Understanding Complexity through Probability*. Harvard University Press, Cambridge, MA.
- . (2005). How are the sciences of complex systems possible? *Philosophy of Science* 72:531–556.
- . (Forthcoming). Probability out of determinism. In C. Beisbart and S. Hartmann (eds.), *Probabilities In Physics*. Oxford University Press, Oxford.