

Toward a Mechanistic Interpretation of Probability

DRAFT

(Please contact me before quoting this version.)

Marshall Abrams

Department of Philosophy

University of Alabama at Birmingham

June 10, 2009

Abstract

I sketch a new objective interpretation of probability, called “mechanistic probability”, and more specifically what I call “far-flung frequency (FFF) mechanistic probability”. FFF mechanistic probability is defined in terms of facts about the causal structure of devices and certain sets of collections of frequencies in the actual world. The relevant kind of causal structure is a generalization of what Strevens (2003) calls microconstancy. Though defined partly in terms of frequencies, FFF mechanistic probability avoids many drawbacks of well-known frequency theories. It at least partly explains stable frequencies, which will usually be close to the values of corresponding mechanistic probabilities; FFF mechanistic probability thus satisfies what in my view is a core desideratum for any objective interpretation. However, FFF mechanistic probabilities are not single case probabilities, and FFF mechanistic probability explains stable frequencies directly rather than by inference from combinations of single case probabilities.

1 Introduction

It's not clear that existing objective interpretations of probability are satisfactory, despite the ubiquity of probability claims in sciences, casinos, and everyday life. This disturbing fact is enough to motivate exploration of the rather complex “mechanistic interpretation of probability” which I'll sketch. Mechanistic probability is partly based on ideas in earlier work of mine and on similar ideas explored in some depth by Michael Strevens (1998; 2003; 2005). As far as I know, Strevens hasn't yet proposed using these ideas as the basis of an interpretation of probability.

2 The machinery: bubbles and microconstancy

2.1 Informal introduction

There are some devices—including many games of chance—whose causal structure is such that it typically matters very little what pattern of inputs the device is given in repeated trials; the pattern of outputs will be roughly the same regardless. For example, a wheel of fortune with red and black wedges—a simplified roulette wheel—is a deterministic device: The angular velocity with which it is spun completely determines whether a red or black outcome will occur. Nevertheless, if the ratio of the size of each red wedge to that of its neighboring black wedge is the same all around the wheel, then over time such a device will generally produce roughly the same frequencies of red and black outcomes, no matter whether a croupier tends to give faster or slower spins of the wheel. Why?

Roughly, the wheel of fortune divides a croupier's spins into small regions (called “bubbles” below) in which the proportion of velocities leading to red and black are approximately the same as in any other such region. As a result, as long as the density curve of a croupier's spins in each bubble is never very steep, *within each bubble* the ratio between numbers of

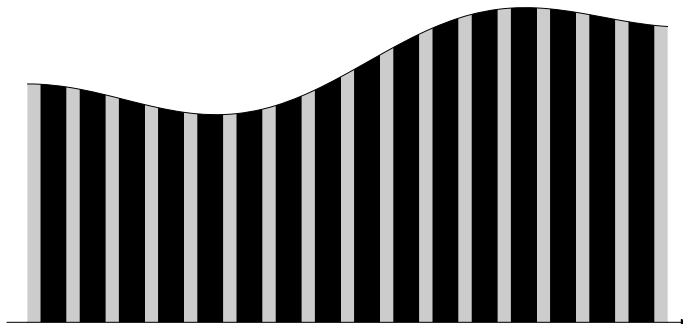


Figure 1: x : velocity; y : spins; gray areas: red frequencies; black areas: black frequencies.

spins leading to red and leading to black will be roughly the same, and will correspond roughly to the sizes of red and black regions within each velocity bubble. The overall ratio between numbers of red and black spins will then be close to the same ratio. In order for this claim to fail, a croupier would have to consistently spin at angular velocities narrowly clustered around a single value, or in a precisely periodic distribution which emphasizes velocities which lead mainly to red (or black) outcomes. Figure 1 illustrates this idea.

2.2 Definitions

Some new technical terms will be helpful. Call a device like the wheel of fortune by which a well-defined set of initial conditions (inputs) produces instances (outputs) of a well-defined algebra of outcomes a *causal map device*; the function from initial conditions to outcomes which it implements is a “causal map”. Outcomes are (perhaps empty) disjunctions of members of a finite partition of *basic outcomes*. An *outcome-inverse set* of initial conditions for a given outcome A is a set of initial conditions all of which would produce the same outcome A ; an outcome-inverse set for A need not include all of the initial conditions which produce A . However, let A^{-1} be that outcome-inverse set for A which does include all of A ’s possible causes. Call an outcome-inverse for a basic outcome a *basic outcome-inverse*. For a causal map device, a *bubble* is a region in the input space containing points leading to all basic outcomes, i.e. which contains outcome-inverse sets for each basic outcome. (As floating

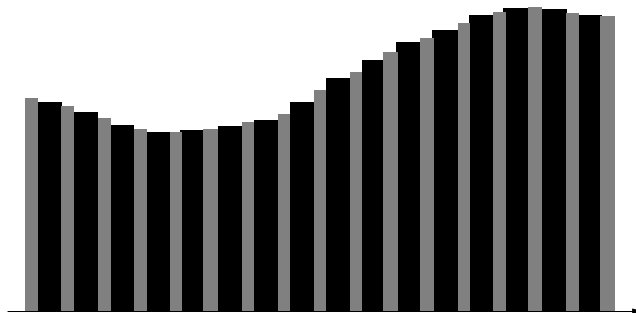


Figure 2: x : velocity; y : spins; gray areas: red frequencies; black areas: black frequencies.

soap bubbles reflect their entire surroundings in miniature, input space bubbles reflect the entire set of outcomes in miniature as well.) A partition of the entire space of possible inputs to the device into bubbles is a *bubble partition*. If there exists a bubble partition with a large number of bubbles, the device is *bubbly*. If each basic outcome-inverse has the same measure conditional on each bubble, i.e. the outcome-inverse’s proportion of each bubble is the same, then the measure is *microconstant* (Strevens, 2003). (I’ll use an upper bar, as in “ \overline{A} ”, to represent negation/complementation.)

In the case of the wheel of fortune, the basic outcomes are red and black, a region of the input space containing velocities leading to the wheel stopping within a pair of red and black regions counts as a bubble, and such a wheel is a bubbly causal map device.

For most croupiers, the wheel of fortune also has a property which Strevens (2003) calls “macroperiodicity”. A causal map device is *macroperiodic* relative to a specification of a density over inputs (e.g. the density curve for our croupier in figure 1), and there is approximately the same proportion of inputs, within each bubble, leading to each outcome.

2.3 A microconstancy inequality theorem

Theorem 1 stated in this section (and proved in appendix A) plays a central role in the rest of the paper. The theorem is similar to central theorems (2.2, 2.3) in (Strevens, 2003), but relaxes some of Strevens’ assumptions.

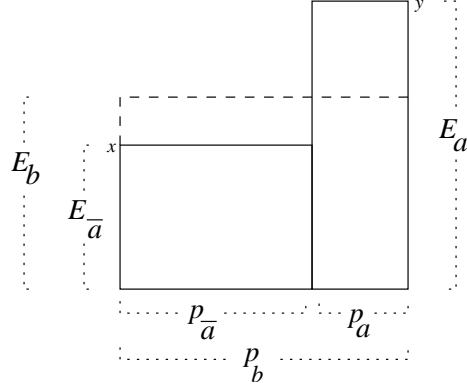


Figure 3: Initial condition distribution within a bubble. See text for explanation.

In the real world there are only a finite number of inputs to a wheel of fortune or any other causal map device. Instead of a smooth density curve over initial conditions, we need only consider a step-shaped function whose height represents the average number of inputs in a given outcome-inverse set within a bubble (figures 2, 3). The number of inputs at any particular point in the input space may be more or less than this value, which is the conditional expectation of numbers of inputs with respect to an (assumed) probability measure over the input space—i.e. conditional on the measure of an outcome-inverse within the bubble. For now we'll focus on a single outcome-inverse set $a = A^{-1} \cap b$ in a single bubble b , along with a 's complement with respect to b , $\bar{a} = b - a = \bar{A}^{-1} \cap b$. In figure 3, the average number of inputs in $A^{-1} \cap b$ is E_a , the height of the right-hand box; the average within $\bar{A}^{-1} \cap b$ is $E_{\bar{a}}$. Also indicated (by the height of the dashed line) is E_b , the number of inputs within b averaged over the entire bubble (the conditional expectation for inputs conditional on the measure of b).

We'll let p_b be the probability of the entire bubble b , p_a be the (unconditional) probability of $A^{-1} \cap b$, and $p_{\bar{a}}$ be the (unconditional) probability of $\bar{A}^{-1} \cap b$ (figure 3). For a wheel of fortune the measure in question might be normalized Lebesgue measure on angular velocities possible for humans. (Though width represents measure in the diagram, the measure need not be a Lebesgue measure.)

Were we considering continuous distributions of inputs, as in figure 1, we could characterize departure from uniformity in terms of the slope of the input distribution. For example, in Poincaré’s (1912, §91) argument that frequencies of red and black outcomes should be roughly equal in a wheel of fortune with equal-sized red and black wedges, Poincaré uses a parameter M which is the maximum slope of a distribution over outcomes. Since I don’t want to assume that the input space is defined by real-valued variables, we need to define an analogue of slope; this could be done in various ways. For example, we can divide the difference between E_a and $E_{\bar{a}}$ by the measure of a given bubble; this value, $(E_a - E_{\bar{a}})/p_b$, which we can call *measure-slope*, is a measure-theoretic analogue of slope. It would correspond to the slope of a line between points x and y in figure 3. A closely related concept will be more useful. Let A^{-1} ’s *bubble-deviation* for bubble b be the difference between E_b and the conditional expected number of inputs for an outcome-inverse within a bubble, divided by the measure of the bubble:

$$\frac{E_b - E_a}{p_b} . \tag{1}$$

Note that this quantity may be negative, as it is for a in figure 3. Measure-slope is related to bubble deviation in that the absolute value of the former is a sum of absolute values of the latter:

$$\left| \frac{E_a - E_{\bar{a}}}{p_b} \right| = \left| \frac{E_b - E_{\bar{a}}}{p_b} \right| + \left| \frac{E_b - E_a}{p_b} \right| . \tag{2}$$

The theorem stated below uses the maximum of absolute values of bubble deviations to constrain the difference between relative frequency of an outcome A and its probability as defined by the measure $\mathbb{P}(A^{-1})$ of inputs which could cause it.

Note that the expectations E_{\bullet} defined above are (conditional) expectations of absolute frequencies; in the end what will matter are relative frequencies, obtainable from absolute frequencies by dividing by N , the total number of inputs in the distribution. Let S be the maximum value for absolute values of bubble-deviations in units of relative frequency; thus

SN is the maximum for bubble-deviations in units of absolute frequencies,

$$S \geq \frac{1}{N} \left| \frac{E_b - E_a}{p_b} \right|, \quad (3)$$

for a given outcome A and for all bubbles b . The theorem (proved in appendix A) can then be stated as follows:

Theorem 1 (microconstancy inequality theorem) *Given a causal map device and a probability measure over its input space which is microconstant with respect to a given bubble partition (it assigns the same outcome-inverse probabilities conditional on each bubble), if S is the maximum of absolute values of bubble-deviations (in units of relative frequency) for A^{-1} , and π is the maximum of measures p_b for all ν bubbles b , then the relative frequency $R(A)$ of A differs from the probability $P(A)$ of A by less than $\nu S\pi^2$, and more precisely by less than $S \sum_b p_b^2$. That is,*

$$\nu S\pi^2 \geq S \sum_b p_b^2 \geq |P(A) - R(A)|. \quad (4)$$

If S is small, the distribution is macroperiodic. The theorem can be applied to multiple outcomes, e.g. to every member of a set of basic outcomes which partition the outcome space:

Corollary 1.1 *If a value S is the maximum bubble-deviation for inverses of all basic outcomes A , the frequencies of all basic outcomes are constrained to be near their probabilities by (4) as well.*

Thus, for a bubbly causal map device with a microconstant input measure, the difference between probability and relative frequency will be small when bubble size (measure) or maximum bubble-deviation (S) are small. Figure 4 illustrates the effect of different bubble sizes with the same relative measures p_a/p_b and $p_{\bar{a}}/p_b$ and the same bubble-deviations and measure slopes. Note how conditional relative frequency of A^{-1} (area of right hand solid box

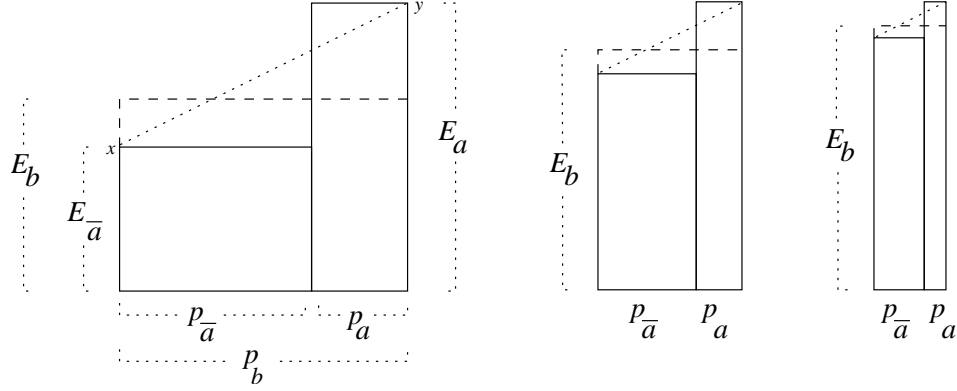


Figure 4: Different p_b 's, # inputs; same conditional probabilities, bubble-deviations, measure-slope.

divided by area in both solid boxes) gets closer to conditional probability as bubble measure is reduced.

More generally, given a bubbly device, for any input measure which is microconstant there will be a large class of input distributions all of which will produce frequencies near to probabilities:

Corollary 1.2 *For any input distributions whose bubble-deviations for an outcome A are less than S_ϵ , the difference between frequency of A and probability of A will be less than $\nu S_\epsilon \pi^2$.*

3 Overview of mechanistic probability

For a bubbly causal map device and a given set of basic outcomes, *if* there is a bubble partition and a probability measure on the input space that gives each outcome the same measure conditional on each bubble, so that the bubble partition is microconstant, then we can define probabilities of outcomes in terms of the measure of initial conditions which lead to them. Then any distribution of inputs with a small bubble-deviations will produce frequencies of outcomes near to their probabilities; the inequality theorem specifies how close

they must be.

Thus if there were some reason that actual distributions of inputs would in fact have small bubble-deviations, we could define a sense of objective probability for outcomes in terms of the following conditions:

- a. The fact that the causal map device is bubbly;
- b. The fact that there's a bubble partition and measure on the input space which determine a small maximum bubble size;
- c. Facts which make it the case that input distributions will have small bubble-deviations relative to this bubble partition and measure.

Outcome probabilities defined by the input measure will then be close to frequencies of outcomes resulting from such distributions.

Such an objective probability would at least partially explain, and in some sense *cause* stable relative frequencies which are near probabilities, since the conditions in terms of which the mechanistic probabilities are defined would guarantee that this is so. Note, however—and this is important—that although the probabilities would be defined by objective conditions, and are defined for a set of basic outcomes of individual trials of the device (where a trial is an actual input to the device causing an outcome to be instantiated), *these need not be probabilities of outcomes on individual trials* in any causal sense. A strict law of nature might determine an outcome of a deterministic trial, and a propensity, if such things exist, might in some sense govern the outcome of an indeterministic trial. Mechanistic probabilities, however, do no such thing. They are objective and help determine frequencies in large numbers of trials; that is all.

What reason is there to assume anything like conditions (b) and (c)? These questions will be the focus of the next section. And what is the point of requiring bubblyness? The answer is that bubblyness makes it *easier* for condition (c) to be satisfied; a wider range of distributions will have small bubble-deviations for a very bubbly device. Alternatively, bubblyness makes

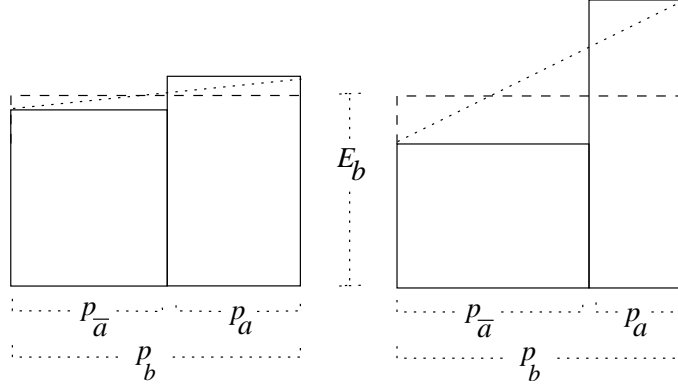


Figure 5: Different p_a 's, $p_{\bar{a}}$'s, bubble-deviations, measure-slopes; same p_b , # of inputs.

it easier for a *measure* to exist which allows a broad set of input distributions to have small bubble-deviations. Figure 5 illustrates this idea, displaying the same number of inputs in A^{-1} and \bar{A}^{-1} with respect to two measures, each of which gives p_b the same value while assigning different pairs of values for $P(A^{-1}|b)$ and $P(\bar{A}^{-1}|b)$.

4 Justification of the input measure

Here I sketch what I think is a reasonable way of defining an input measure for a causal map device, in terms of what I call “far-flung frequency” (FFF). This will provide the basis of what I call “FFF mechanistic probability”. (There are ideas in Strevens’ book (2003) which might justify an input measure but I don’t think they are widely applicable, as I argue in a longer version of this paper.)

Given a set of basic outcomes for a causal map device which is bubbly with respect to those outcomes, we need a general way of specifying a probability measure on the device’s input space that allows the possibility (a) that the measure is microconstant—i.e. that the measure of a basic outcome-inverse conditional on a bubble is the same for every bubble—and (b) that the measure in some sense reflects patterns of inputs in the world in such a way that relevant collections of actual inputs to such devices tend to be macroperiodic—i.e.

have small bubble-deviations. Our method of specifying an input measure should therefore reflect in some way facts about actual inputs to the device in question and similar devices.

An unhelpful method of defining an input measure for an actual causal map device (e.g. a wheel of fortune) during an interval of time T would be to define the input measure as equal to frequencies of actual inputs to the device during that period of time. This would be to define mechanistic probability as simple actual frequency. However, we can cast our net more broadly and include inputs to a (token) causal map device D and similar devices both before and after T . How similar? How much before and after?

The actual inputs to devices that matter for defining the input measure for a device D should at a minimum be inputs to devices with the roughly same input space as D ; they need not map inputs to outcomes in the same way as D , however. As a first pass, we might define the input space of an actual wheel of fortune D in some way which reflects patterns of inputs to all actual devices with roughly the same input space of angular velocities over a large period of time beginning in the past and going into the future. For example, we should at least include angular velocities of wheels produced by actual croupiers for wheels of fortune and roulette wheels for the last 500 years and next 500 years. We need to maintain the distinction between collections of actual inputs by different croupiers, and then define D 's input measure in terms of each of these separate collections of inputs. This move, along with procedures outlined below, will allow us to distinguish between our world and some worlds where every croupier simply spins wheels at his or her own special narrow region of velocities.

Define a *natural collection of inputs* with respect to a causal map device D to be:

A large set of all and only those actual inputs, produced by a single device E to a causal map device D^* with the same input space as D (and perhaps other loose similarities) during a single interval of time T .

The producing device E , for example, might be a human croupier. We'll end up defining the

input measure for a causal map device in terms of relative frequencies in natural collections of inputs.

The fact that a natural collection must include all inputs during an interval T means that the collection can't just include, say, the inputs which produce red outcomes during T . The fact that a natural collection must be produced by a single device, and is not just the "supercollection" of all inputs produced by by all such devices (e.g. all croupiers) avoids other problems. For example, using such a supercollection to define the input measure roughly as we do below would mean that we couldn't distinguish between the actual world, in which frequencies on wheels of fortune (roulette wheels, etc.) are often close to probabilities, and a world in which inputs to wheel D_1 always remain in small interval x_1 of angular velocities, inputs to wheel D_2 always remain within a different small interval x_2 , and so on for all relevantly similar devices (e.g. because all croupiers where very heavy weights on their wrists). Collect all of those disparate sets of inputs together into one collection, and you could have exactly the same overall frequencies as in the actual world; this would end up determining the same mechanistic probabilities for outcomes of our device D . But no such device would have frequencies of outcomes close to probabilities. It's not an acceptable implication of an interpretation of probability that the world could routinely make all frequencies of a given outcome be far from its probability.

This specification of what counts as a natural collection of inputs for a given causal map device is quite loose. One question is how far the spatial and temporal boundaries across which these collections are to be found should extend; I address this question below.

Since this approach uses frequencies in these natural collections of inputs to similar devices over a wide spatiotemporal region to define the input measure which will define outcome probabilities, I call the resulting kind of mechanistic probability *far-flung frequency (FFF) mechanistic probability*.

Given such a set of natural collections of inputs, we define the input space of a bubbly

causal map device D by choosing an input probability measure that minimizes the differences between relative frequencies in each collection and input probabilities. More precisely, my proposal is that we choose measures for bubbles and for outcome-inverses which minimize the sum of squares of bubble-deviations for all of the relevant natural collections.

Thus let c index natural collections, b index bubbles, and a index outcome-inverses. p_b is input probability of bubble b , and p_a is the input probability of outcome-inverse a , which we require to be the same conditional on every bubble. Similar conventions will apply to expected numbers of inputs E_\bullet , defined above; e.g. E_{cba} is the expected number of inputs from natural collection c within outcome-inverse a in bubble b . Let N_c be the size of natural collection c and w_c be a weight function which depends on N_c . Then the quantity to be minimized by adjusting the p_b 's and p_a 's is:

$$\sum_c \sum_b \sum_a w_c \left(\frac{E_{cb} - E_{cba}}{p_b} \right)^2$$

The quantity squared is the bubble-deviation, which is the starting point of the proof of the inequality theorem. We should give additional weight to the contribution of bubble-deviations in larger natural collections; this formula does this partly by the fact that it uses bubble-deviations in units of raw numbers of events rather than relative frequencies. w_c is an additional weighting function which can be used to adjust the weighting; I'm exploring possible definitions of w_c . The entire expression in effect measures the distance between a set of averages with respect to bubbles E_{cb} and a set of averages with respect to measure of frequencies over outcome-inverses in bubbles, weighted in particular ways by collection sizes and bubble measures. (I'm currently investigating the properties of this function, including the conditions under which the resulting measure is unique.)

There is no guarantee, though, the input measure so defined actually does produce small bubble-deviations (macroperiodicity) for most natural collections; the natural collections might just vary too wildly. However, if the device is bubbly, and if the minimization produces an input measure which makes bubbles small, then it will be relatively easy for a

natural collection to produce small bubble-deviations. If most of the natural collections are *not* macroperiodic with respect to the device’s input measure, that means that the kind of processes that produce inputs to such devices are simply too varied in their effects even for a bubbly device to generate mechanistic probability. In such a case we have some of the ingredients for mechanistic probability (a bubbly causal map device), and we have constructed a good input measure, but it’s not good enough, and there is thus no mechanistic probability in this instance. We might require, for example, that maximum bubble deviations be no more than .05 in 95% of the natural collections.

Now I am ready to answer the question of how extensive the set of natural collections should be. If expanding or contracting the size of the spatiotemporal region in terms of which natural collections are defined makes a significant difference, then mechanistic probability does not exist. For example, suppose that members of a set of natural collections are for the most part macroperiodic relative to the constructed input measure for a bubbly causal map device—and so, *prima facie*, mechanistic probability exists in this case. Suppose in addition that expanding the region a bit (e.g. by another 1000 years in either direction), constructing the input measure with the new set of collections, and then testing for their general macroperiodicity, would not satisfy the requirements for mechanistic probability. Then we do not have an instance of mechanistic probability. This is a kind of “If you have to ask, you don’t belong” rule.

5 FFF mechanistic probability

FFF Mechanistic probability in one or another sense exists for a particular actual causal map device and a specified partition of its output space into basic outcomes if

1. That the causal map device is bubbly;
2. There are many large natural collections of inputs to “relevantly” similar devices;

3. The microconstant input measure constructed through a minimization procedure (as above) from these natural collections makes most of the collections macroperiodic relative to this input space and the device's bubble partition.
4. Expanding or shrinking the spatial or temporal range across which natural collections are defined a bit would not make much difference to the preceding.

Then the mechanistic probability of an outcome is defined by the input probability of its inverse.

Note that most devices which help constitute mechanistic probability will produce outcome frequencies which are close to outcome probabilities. It is largely bubblieness that forces frequencies of outcomes into stable patterns. Bubblieness as it were “compresses” the wildness of frequencies in collections of initial conditions into the regularity of outcome frequencies. So mechanistic probability largely explains stable frequencies.

We can go further. If manipulability is the mark of causation (Woodward, 2003), then it appears that mechanistic probability causes outcome frequencies. For we can manipulate frequencies of outcomes by modifying the structure of the device so that larger or smaller sets of initial conditions are mapped to various outcomes—as we do when we manipulate the size of wedges on the wheel of fortune or roulette wheel, or as we do when we load a pair of dice. Such manipulations do not change the input measure, nor do they change frequencies of initial conditions in the world. The stability of frequencies is explained by the structure of the mechanism and the relative lack of diversity in natural collections (i.e. lack of diversity relative to the particular bubbly structure of the device). The particular frequencies which are stable are determined by the structure of the mechanism, which is part of what constitutes mechanistic probabilities of outcomes.

Furthermore, nearby worlds in which natural collections differ a bit will be ones in which frequencies of outcomes will generally be similar to what they are in the actual world. This is one main reason why FFF mechanistic probability is superior to a simple actual finite

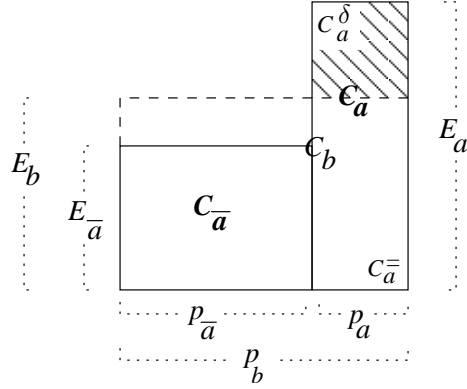


Figure 6: Initial condition distribution within a bubble. See text for explanation.

frequency theory of probability. More generally, bubblieness means that the fact that the relevant set of natural collections is defined in an (admittedly) vague way doesn't matter much. Because bubblieness, microconstancy, and general macroperiodicity buffers outcome frequencies against not-too-radical differences in the set of natural collections means that mechanistic probability is still a sensible notion even though part of its basis is defined loosely. (Note however that what constitutes mechanistic probability are actual facts—not counterfactuals—concerning the causal structure of the device and collections of actual events. It is actual conditions and events which are truthmakers for FFF mechanistic probability.)

Finally note that mechanistic probabilities are not single case probabilities, though they are objective probabilities which causally explain frequencies. Mechanistic probability explains relative frequency directly rather than by derivation from single case probabilities for multiple trials.

6 Appendices

A Proof of microconstancy inequality theorem

Figure 6 is an enhanced version of figure 3. See Section 2.3 for definitions of concepts and most symbols. In addition, the areas C_a and $C_{\bar{a}}$ of the two solid boxes in figure 6 represent the absolute frequencies (counts) of inputs in each of the regions $A^{-1} \cap b$ and $\bar{A}^{-1} \cap b$ respectively. Note that $C_a = p_a E_a$, i.e. $E_a = C_a/p_a$; $C_{\bar{a}}$ and $E_{\bar{a}}$ are similarly related. C_b , the total number of inputs to b , is represented by the area under the dashed line and is equal to $p_b E_b$. ($C_a^=$ and C_a^δ will be defined below.)

By the definition of S , i.e. (3), we have $SNp_b \geq |E_b - E_a|$, and

$$SNp_b^2 \geq p_a |E_b - E_a|, \quad (5)$$

since $p_b \geq p_a$. (All quantities which are not defined in terms of subtraction above are required to be positive and so can be freely moved in and out of the absolute value delimiters.)

Frequency and measure would be equal if the distribution were constant over each bubble, as Strevens' theorem shows. However, whether or not the distribution is constant over a given bubble b , we can compare the actual number of inputs in a with the number of inputs that would have to obtain in a if relative frequencies within b did equal relative measures of p_a and $p_{\bar{a}}$. That is, we compare the number C_a of actual inputs to a with those that would occur if conditional expectation over both a and \bar{a} were equal to the actual conditional expectation E_b for the entire bubble. Thus we define the quantities $C_a^= = p_a E_b$ and $C_a^\delta = C_a^= - C_a$. (In figure 6, C_a^δ is negative, indicated by shading; the corresponding value for $C_{\bar{a}}$ would be positive.) Thus $C_a^\delta = p_a E_b - p_a E_a$, so from (5) we have

$$SNp_b^2 \geq |C_a^= - C_a| = \left| C_b \frac{C_a^=}{C_b} - C_a \right|.$$

However, as figure 6 illustrates, $C_a^=/C_b$ is equal to p_a/p_b since $C_a^= = p_a E_b$ and $C_b = p_b E_b$.

Note that $p_a/p_b = \mathbf{P}(A^{-1}|b)$ since $p_a = \mathbf{P}(A^{-1} \cap b)$, so

$$SNp_b^2 \geq |C_b\mathbf{P}(A^{-1}|b) - C_a| .$$

If we now sum over all bubbles b , we have

$$SN \sum_b p_b^2 \geq \sum_b |C_b\mathbf{P}(A^{-1}|b) - C_a| \geq \left| \sum_b C_b\mathbf{P}(A^{-1}|b) - \sum_b C_a \right| .$$

On the assumption that the bubble partition is microconstant with respect to the input measure, the probability of A^{-1} conditional on every bubble is the same and is equal to $\mathbf{P}(A^{-1})$, so since $C_a = C_{A^{-1} \cap b}$,

$$SN \sum_b p_b^2 \geq \left| \mathbf{P}(A^{-1}) \sum_b C_b - \sum_b C_{A^{-1} \cap b} \right| .$$

But $\sum_b C_b$ is just the sum of numbers of inputs to all bubbles, i.e. the total number of inputs of any kind, N ; and $\sum_b C_{A^{-1} \cap b}$ is the total number of inputs in A^{-1} across the entire input space. Dividing by N , we get

$$S \sum_b p_b^2 \geq |\mathbf{P}(A^{-1}) - \mathbf{R}(A^{-1})| = |\mathbf{P}(A) - \mathbf{R}(A)| , \quad (6)$$

where $\mathbf{R}(A^{-1})$ is the relative frequency of inputs leading to A , and the probability $\mathbf{P}(A)$ and relative frequency $\mathbf{R}(A)$ of A are simply the probability and relative frequency of inputs which cause instances of A . Finally, if π is the maximum of bubbles' measures (cf. Strevens' "constant ratio index") and ν is the number of bubbles,

$$\nu S \pi^2 = S \sum_b \pi^2 \geq |\mathbf{P}(A) - \mathbf{R}(A)| \quad (7)$$

since $\sum_b \pi^2$ just adds ν instances of π^2 . (If all bubbles have the same measure, the quantity on the left is equal to $S\pi$, since then π is the measure of every bubble, and therefore $\nu\pi = 1$.)

From (6) and (7) we have:

Theorem 1 (microconstancy inequality theorem) *Given a causal map device and a probability measure over its input space which is microconstant with respect to a given bubble*

partition, if S is the maximum of absolute values of bubble-deviations (in units of relative frequency) for A^{-1} , and π is the maximum of measures p_b for all ν bubbles b , then the relative frequency $R(A)$ of A differs from the probability $P(A)$ of A by less than $\nu S\pi^2$, and more precisely by less than $S\sum_b p_b^2$. That is,

$$\nu S\pi^2 \geq S\sum_b p_b^2 \geq |P(A) - R(A)|.$$

B Strevens' terminology

For a reader familiar with Strevens' work here are some relationships between those of my terms not taken from Strevens and some of his terms. Much of Strevens' extensive technical vocabulary is not optimal for my purposes.

In Strevens' (2003) terms, what I call a "causal map device" is, roughly, a "mechanism" provide with a "designated set of outcomes"; if we specify in addition a distribution over initial conditions, we have what Strevens calls a "probabilistic experiment". "Bubble partition" is a generalization of Strevens' notion of a constant ratio partition (CRP), and "bubble" is a generalization of "member of a constant ratio partition". A bubble partition is a CRP (and is therefore microconstant) if we consider only two possible outcomes A and \bar{A} , and if the measure of A^{-1} conditional on each bubble is the same, i.e. the same proportion of each bubble consists of causes of A rather than \bar{A} .

Note that macroperiodicity concerns proportions of inputs within each bubble, whereas the definitions of "constant ratio partition" and "microconstant" are in terms of an underlying measure which need not equal relative frequencies of inputs.

References

Poincaré, H. (1912). *Calcul des Probabilités* (2nd ed.). Gauthier-Villars. (1st ed. 1896).

Strevens, M. (1998). Inferring probabilities from symmetries. *Noûs* 32, 231–246.

Strevens, M. (2003). *Bigger Than Chaos: Understanding Complexity through Probability*.
Harvard University Press.

Strevens, M. (2005). How are the sciences of complex systems possible? *Philosophy of Science* 72(4), 531–556.

Woodward, J. (2003). *Making Things Happen*. Oxford University Press.