

# Humean Effective Strategies

Carl Hoefer

*ICREA & Universidad Aut3noma de Barcelona*

Draft November 2003

**ABSTRACT:** In a now-classic paper, Nancy Cartwright argued that the Humean conception of causation as mere regular co-occurrence is too weak to make sense of our everyday and scientific practices. Specifically she claimed that in order to understand our reasoning about, and uses of, effective strategies, we need a metaphysically stronger notion of causation and causal laws than Humeanism allows.

Cartwright's arguments were formulated in the framework of probabilistic causation, and it is precisely in the domain of (objective) probabilities that I am interested in defending a form of Humeanism. In this paper I will unpack some examples of effective strategies and discuss how well they fit the framework of causal laws and criteria such as *CC* from Cartwright's and others' works on probabilistic causality. As part of this discussion, I will also consider the concept or concepts of objective probability presupposed in these works. I will argue that Cartwright's notion of a *nomological machine*, or a *mechanism* as defined by Stuart Glennan, is better suited for making sense of effective strategies, and therefore that a metaphysically primitive notion of causal law (or singular causation, or capacity, as Cartwright argues in (1989)) is not – here, at least – needed. These conclusions, as well as the concept of objective probabilities I defend, are largely in harmony with claims Cartwright defends in *The Dappled World*. My discussion aims, thus, to bring out into the open how far Cartwright's current views are from a radically anti-Humean, causal-fundamentalist picture.

**0. Introduction.** Throughout her career, Nancy Cartwright has consistently argued against the Humean prejudices of her logical empiricist predecessors, at least in the areas of causality and the epistemology of science. The first assault in her campaign was the classic paper "Causal Laws and Effective Strategies" (1979, 1983 ch. 1). This paper argues for two main theses. First, that there is no way to reduce facts about causation to facts about probabilistic relations; and second, that in order to understand the effective strategies we use to achieve desired results, we need to invoke a strong notion of *causal laws*. When we know that it is a causal law that C brings about E (or raises the level of E, or makes E more probable, . . .), then we have an effective strategy for E. It is only the second thesis that I will be attempting to undermine in this paper, by showing that the talk of causal laws, and the implicit picture of the (1979) paper, have some serious faults.

What I will do is focus on aspects of the problematic that Cartwright glosses over relatively briefly, and try to show how a slightly different way of thinking about things can work equally well – perhaps better – at uncovering and describing our effective strategies. The point will be to show that this different perspective is wholly compatible with *Humeanism about (real, or objective) probabilities*, and with *agnosticism about causation* as a primitive relation (i.e., causal agnosticism as opposed to causal fundamentalism). The goal is to show how we can account for our effective strategies, without buying wholesale into an ontology of causal laws, singular causation,

and capacities. In this sense, I will defend the spirit of Humeanism about causation, at least in a small way. I will not, however, try to argue for a view effective strategies that is purged of any taint of causal talk.

In elaborating a different view of effective strategies, I will borrow heavily from some leading ideas of Cartwright's latest book, *The Dappled World*.<sup>1</sup> This paper is therefore an attempt both to defend much of the perspective offered in *The Dappled World*, and to show that it contrasts strongly with the causal-fundamentalist picture to be found in some of Cartwright's earlier works.

**1. Cartwright against Humean probabilistic causation.** In this section I will describe the main points of Cartwright's (1979) paper, including the famous criterion *CC* and three key examples of effective strategies (Malaria, TIAA-CREF, heart disease). Cartwright uses her examples to argue very convincingly that a Humean account of causation that seeks to reduce causal facts to facts about probability relations is doomed to fail. What she offers in its stead is a species of what might be called "causal fundamentalism",<sup>2</sup> namely a view that takes "causal laws" to be fundamental facts of our universe.

"If indeed, it *isn't true* that buying a TIAA policy is an effective way to lengthen one's life, but stopping smoking is, the difference between the two depends on the causal laws of our universe, and on nothing weaker." (1983), p. 22.

I have never felt I understand what a "causal law" is, and in (1983) Cartwright does not give us an explicit definition. We do however get an implicit definition: At least, the true statements " $C \rightarrow E$ " that pass the test of principle *CC* should be counted as causal laws. Later we will come back to the issue of what constitutes a causal law.

The basic idea of a Humean reductive theory of (probabilistic) causation is that  $C$  causes  $E$  if  $P(E|C) > P(E|\neg C)$  and some other conditions (all of which should be non-causal, i.e. compatible with whatever Humeanism is in play) are satisfied as well.  $C$  and  $E$  should be event types, not particulars, at least in the class of theories of interest to us here. Since there are many well-known examples of factors that increase the probability of an effect  $E$  *via* spurious correlation rather than causation, the real content of the theory will naturally be in the extra conditions. For example, since the probability of lightning is greatly increased by the presence, less than ten seconds later, of thunder, we need a condition that helps our theory rule out thunder being a cause of lightning. Suppes' (1970) theory of probabilistic causation, one of the first, took a sensible approach to cases like this: insist that the cause  $C$  must occur before the effect  $E$ . This by no means finishes the task of eliminating spurious correlations, unfortunately. There are ubiquitous cases of effects of a common cause ( $D \rightarrow C$  and  $D \rightarrow E$ ), where  $C$  regularly happens before  $E$  and is strongly positively correlated with  $E$ ,

---

<sup>1</sup> After I presented the main contents of this article in Oviedo (LMPS '03), Paul Teller pointed out to me that Stuart Glennan has written articles defending a mechanism-based view of causality that is very close to some of the ideas I advocate here. See Glennan (1996, 1997, 2002).

<sup>2</sup> Turnabout is fair play: Cartwright's philosophical opponents who believe in fundamental laws of nature may deserve the epithet "fundamentalist", but she often seems to be no less a fundamentalist about causation. John Norton (2003) used this term first, I think, and I gladly borrow it from him.

but is not a cause of E. Ruling out cases of this form, and a variety of more complicated forms, is a job that has never proved possible, at least within the strictures of Humeanism.

Cartwright offers cases with the probability structure known as Simpson's paradox to illustrate her general argument against the Humean approach to probabilistic causation. Let's take the smoking/heart disease thought example (the probabilities we will posit are by no means true of any actual populations). Suppose that it were found that, in the statistics for the whole adult population,  $P(\text{HD}|\text{SM}) < P(\text{HD}|\neg\text{SM})$ . This could happen even if smoking is in fact a cause of heart disease and not a preventer of it. How? Well, suppose that regular exercising is a strong preventer of heart disease, and that as it happens, the frequency of regular exercising is much higher in the smoking population than in the non-smoking population. Then the probability relation mentioned above could hold, yet when we partition the population into exercisers and non-exercisers (and "hold fixed" this factor, conditionalize on it), the probabilistic significance of smoking reverses:  $P(\text{HD}|\text{SM} \ \& \ \text{EX}) > P(\text{HD}|\neg\text{SM} \ \& \ \text{EX})$  and  $P(\text{HD}|\text{SM} \ \& \ \neg\text{EX}) > P(\text{HD}|\neg\text{SM} \ \& \ \neg\text{EX})$ . And these probabilities, we are to take it, reflect the true causal facts, that smoking does cause heart disease.

The point Cartwright makes with these examples is simple but devastating: if there are other causal factors relevant to an effect E (positively or negatively) that may induce misleading probabilities, we have to hold them fixed in order for the probability of E given C to genuinely reflect the fact that C [causes/prevents] E. So if there are five other genuine causes/preventers of a given E,  $C_i$  for  $i = 1$  to 5, then in order to judge whether C causes E what we need to look at is the probabilistic relevance of C for E, in each of the subpopulations where each of these five factors  $C_i$  is held fixed (positively or negatively), e.g.,  $P(E|C \ \& \ C_1 \ \& \ \neg C_2 \ \& \ C_3 \ \& \ C_4 \ \& \ \neg C_5)$ . Formalizing this notion Cartwright gets CC:

"CC:  $C \rightarrow E$  iff  $\text{Prob}(E|C \ \& \ K_j) > \text{Prob}(E|K_j)$  for all state descriptions  $K_j$  over the set  $\{C_i\}$ , where  $\{C_i\}$  satisfies  
 (i)  $C_i \in \{C_i\} \Rightarrow C_i \rightarrow +/- E$   
 (ii)  $C \notin \{C_i\}$   
 (iii)  $\forall D (D \rightarrow +/- E \Rightarrow D = C \text{ or } D \in \{C_i\})$   
 (iv)  $C_i \in \{C_i\} \Rightarrow \neg (C \rightarrow C_i)$ ."<sup>3</sup>

This is not, of course, an analysis or definition of the causal relation  $\rightarrow$  in terms of probability, because the relation occurs on both sides of the *iff*. It is, rather, as Cartwright puts it, ". . . the strongest connection that can be drawn between causal laws and laws of association."<sup>4</sup>

It is also a disaster for the basic Humean programme of reducing causation to probabilistic relations. Viewing it for the moment as an epistemic recipe, what CC says is that in order to infer that C is a cause of E from probabilities, one has to first know *all the other causes* of E, and examine the effect of C on E in *each* of the subpopulations holding fixed a combination +/- of these other causal factors. CC's truth (if it is true) does not logically preclude a successful reduction of causation to probabilistic facts. But it does make it look rather unlikely, and makes it more natural to see the logical relationship going in the other direction: causal facts are (logically, or ontologically)

<sup>3</sup>(1983), p. 26. Condition (iv) is needed to handle problems that would occur if one held fixed causes of E that sometimes are intermediate steps between C and E. CC is not itself immune to counterexamples and problems; see for example Otte (1985).

<sup>4</sup>(1983), p. 26.

prior, and give rise to the probabilistic facts. And that is part of Cartwright's causal fundamentalist view: causal relations *give rise to* probabilistic relations by their operation, and the latter are at best a dubious tool to be used in trying to infer the existence of the former.

So much for causal laws, as implicitly defined by *CC*. The application to effective strategies is straightforward: If  $C \rightarrow E$ , then introducing (or augmenting, increasing, . . .) *C* is an effective strategy for bringing about *E*, in all circumstances. Joining TIAA is not an effective strategy for extending ones' life, because it is (presumably) no causal law that joining TIAA  $\rightarrow$  longer life. And spraying oil on swamps is an effective strategy for preventing malaria, while burning the blankets of the sick is not, because of the (presence/absence) of the corresponding causal laws linking these event types.

Before we take a critical look at the examples of causal laws and effective strategies used by Cartwright, we need to pause for a moment to think about the *objective probabilities* being used in the discussion.<sup>5</sup> She discusses the question in section 2.2 of *CLES*, and insists that they must be understood as simply sufficiently-stable [actual] frequencies. She does not want to make a stronger linkage between probability and causation possible by going metaphysical, opting for some primitive notion of propensity or a translation into counterfactuals. And with this I am in full agreement: there is no call to ruin a perfectly good notion like objective probability, just because we can't make it link up nicely with facts about causation.

“Probabilities serve many other concerns than causal reasoning and it is best to keep the two as separate as possible. In his *Grammar of Science* Karl Pearson taught that probabilities should be theory free, and I agree.” (1983), p. 39.

As we will see in section 3, however, this simple frequentist view of objective probabilities lands *CC* and the associated view of causal laws in great difficulties. But before we get to these, in the next section I want to lay out the elements of a different – but equally empiricist, equally Humean – account of objective chance.

## 2. Humean objective chance.

The need for a Humean account of objective probabilities (or chance, as I will usually say) different from simple actual frequentism is not hard to see. Though not all – perhaps not even most – of the traditional complaints against actual frequentism are sound<sup>6</sup>, still there are some glaring problems that have made the view a nearly extinct species in recent decades. First of all, we expect the actual frequencies of things to at best come close to the real probabilities, an expectation that is weaker the smaller the number of actual cases involved. For example, the proportion of heads among well-flipped coins, in the history of the world, is no doubt near 0.5, but it is also no doubt *not exactly* 0.5. Yet it would be nice to have a way to say that 0.5 is in fact the correct value of the probability. The example can be strengthened by considering similar cases where the numbers are much smaller. Suppose that a proper roulette wheel with exactly 25 slots was only built once in history, and used just briefly in an obscure French

<sup>5</sup>Cartwright is explicit that these probabilities must be objective, not subjective, and indeed the reason is obvious: my (or anyone, including an “ideal rational agent”) having certain degrees of belief cannot make it the case that *C* causes *E*, nor that *C* is an effective strategy for bringing about *E*.

<sup>6</sup>For a compendium of these arguments, see Hajek (199?).

casino. And suppose that the ball only fell into the 00 slot on that wheel in 2.333 percent of the spins. It would nevertheless be nice to be able to say that the probability was in fact 4%, without going off the metaphysical deep end in order to do so.

A second problem, perhaps worse, is that statistics and frequencies are ubiquitous, but not all of them should be thought of as probabilities. The frequency of men with only silver coins in their pockets on Tuesdays that are their birthdays, in the whole population of such men (on such days) is a statistic of no meaning and no utility. Notice that this sort of statistic does not support a temporal reading such as to potentially guide expectations. It is *not* to be identified with the probability that, on my next Tuesday birthday, I will end up having only silver coins in my pocket. It can be converted to a probability if we instead read it as the probability of *getting someone who only has silver coins in his pocket*, if one *randomly selects a man on one of his Tuesday birthdays* from the entire pool of such individuals over all history. The latter sort of gloss allows one to turn any mere statistic into a genuine objective probability, but this does not justify the former sort of reading, which is what we mostly would like to have.<sup>7</sup>

It is commonplace now to insist, following Ian Hacking, that objective probabilities can only be associated to proper *chance setups*. Frequentism does not build in this restriction, even if we add a requirement of stability. (The method of translating any statistic into a probability in footnote 7 is, in effect, a method of building the statistic into a proper chance setup.) But what sort of a thing is a chance setup, and how should we motivate the distinction between proper objective chances and mere statistics?

The account I will briefly sketch reflects the trajectory of my own interest in probabilities, which began with an interest in David Lewis' (1994) account and grew off in a different direction from there. Cartwright has never had much sympathy for Lewis-style Humean programmes, and her interest in objective probabilities has always been closer to the needs and practices of ordinary science. Despite this difference, I think the account I will sketch here is very close to the account developed in Cartwright (1999), chapter 7.<sup>8</sup>

**2.1 What Chances are For.** If you know that something is the case, or you know that some other thing is definitely going to happen, then you are all set; knowing the probability of those things is then at most of academic interest to you. But often we have to work in circumstances of ignorance. I don't know whether it will rain tomorrow, so knowing the objective probability that it will (if such a thing exists) would be very useful to me. If it is less than 20% I will wear my new shoes and not take an umbrella, but if it is more than 80% I will dress warmly, wear old shoes, and carry my

---

<sup>7</sup>Generally when we use a statistic, we would like to give it a sort of causal (or perhaps better, expectation-guiding) reading. For example, we would like to use the statistics concerning incidence of breast cancer in a certain population as though it gave us the probability that a person who is about to enter that population group contracts breast cancer while a member of the group. But it is no such thing (at least, on the face of things). It is only a genuine objective probability if read as the probability of obtaining a person who has breast cancer, if one randomly samples one person from the population. And *this* objective probability is, unfortunately, rarely of use or interest to us.

<sup>8</sup>See Hoefer (1997), (2003). My views have developed mainly out of a desire to correct and perfect the Lewisian approach to chance, but have certainly been influenced also by reading Cartwright's works and discussing many issues with her, during the years 1998 - 2002.

big umbrella. Objective probability is, in the now well-known phrase, a “guide to life”. That is its nature or essence, if you like, and this role is neatly captured in David Lewis’ “Principal Principle”, which says roughly:

*PP*: If you have no background knowledge relevant to whether or not it is (or will be) the case that *A*, other than perhaps background knowledge concerning the objective chance of *A*, then if you come to believe that the objective chance of *A* is in fact *x*, your subjective degree of belief in the truth of (or coming to pass that) *A*, should also be *x*.<sup>9</sup>

PP is meant to be a rather obvious principle of rationality: if you don’t follow it, you are either being perverse in some way, or falling short of logical coherence, or you simply don’t understand the concept of objective chance. *The point of saying that* the probability of 6 upon rolling a fair die is 1/6 is precisely to indicate what a rational degree of belief (hence rational/fair betting behavior, etc.) in that outcome is. It is not just a shorthand way of saying what the actual frequency of 6’s is, in the past or even in all of history, though we do expect chances and frequencies to be numerically close, in most cases. Nor is it a way of saying that there are 6 possible ways for a die to land and that we are indifferent between them. Not only are there other cases where we can “be indifferent” in two or more ways, yielding contradictory prescriptions for the probabilities, but moreover mere indifference is no grounds for saying what objective chances are. If you know *nothing* about a die that someone hands you, then you certainly don’t know what its chance of landing 6 is! On the other hand if you know that it is a perfect cube (with rounded edges), has uniform density and is not magnetic, etc., then you may indeed have grounds for saying that the chance of heads is 1/6, but these grounds are not best thought of as a matter of “indifference”.

Finally, to say that the chance of heads is 1/6 is not to attribute a mysterious causal power to the die that “necessitates” a roll of 6 – but only to the degree 1/6. Whatever that might mean. There are too many varieties of propensity theories of chance to try to survey them here, but what I want to emphasize is that whatever objective chances are, they are certainly compatible with the reign of determinism at the level of physical law (contrary to what at least many propensity theorists claim).<sup>10</sup> One of the claims I will argue for below is that there may be fewer objective chances out there than some people assume. But I would argue strenuously that it can’t be the case that there are *none* (or none whose value is neither zero nor one). Objective probabilities are the kinds of features of reality displayed *par excellence* in gambling devices and coin flippings, and presumably radium decays and many other phenomena.<sup>11</sup> A view which says that there are no objective chances if the world turns out to be at bottom deterministic, is in my view just changing the topic of conversation. Even if the world is deterministic, we (in our ignorance) still need all the guides to life we can get. There are indeed features of reality that we can see will serve to play the

---

<sup>9</sup>I eschew the usual mathematical formulations of *PP* here in order to make its common sense nature more clear.

<sup>10</sup>And also contrary to David Lewis. See Hoefer (2003) for discussion of why linking objective chance to indeterminism is a mistake.

<sup>11</sup>But not necessarily all the phenomena that we typically pretend have objective chances. For example, it is far from clear to me that there is an objective chance of it raining tomorrow, in Castelldefels (Spain). (In Europe, unlike the U.S., weather forecasters rarely give numerical probabilities in their forecasts.)

role of guide as specified in *PP* – such as the features of a fair die mentioned earlier, plus what we know about how people throw dice and how they bounce, etc. – so there are objective probabilities in the world

Like David Lewis, I claim that *what chances are for* (as expressed in *PP*) is our best guide to what chances *are*. Objective chances must, at the least, be facts that entail the rationality or correctness, in some sense, of the Principal Principle. Now I will sketch a Humean view of objective chance that is meant to satisfy this constraint.

**2.2 What chances are.** A proper Humean empiricist will insist that objective probabilities, whatever they may be, must at least supervene on the sum total of actual events in world history. They are not some mysterious or hidden springs lurking underneath (as some views take the laws of nature to be) and forcing the world's events to be the way they are. Instead they are patterns that can be discerned in the vast panoply of events occurring in the world. What kind of patterns? Finite frequentism answers the question in a simple way: relative frequencies. Or perhaps: relative frequencies meeting certain tests of stability and distribution. But there are too many of these relative frequencies, and that undermines the sensibility of *PP*. The chance of rain tomorrow in Castelldefels should be defined as the relative frequency of rain-the-next-day in a reference class of preceding-days “like today”. But – like today, in which respects? If we specify too many respects, we whittle our reference class down to nothing, or nearly-nothing, in which case it would seem wrong to let the frequency guide our credence. (If the chance of rain does exist, I am certain that it is neither 0 nor 1.0!)<sup>12</sup> On the other hand, there may be no good reason (from the perspective of simple Humean frequentism) to choose one set of attributes that days “like today” share, over another set; and the other set will likely give different frequencies. This is why it is better to let go of frequentism,<sup>13</sup> and move to a more sophisticated Humean account based on the two key notions of *best systems* and *nomological machines*.

Lewis (1994) offers a package account of laws and objective chances together, one that in effect says objective chances (if they exist) are dictated by laws of nature. It is called a *best systems* account because it meets the demands of Humeanism by defining the laws of nature as a set of axioms that systematize the patterns in actual occurrent events, obtaining a “best” combination of simplicity and strength. In our world, it may be that the best system of axioms we can have does not deterministically specify what *will be* the case, always and everywhere, but rather tells us the objective probability of various occurrences. These then are the objective chances.

---

<sup>12</sup>Here I am presupposing that the frequencies in *past cases* determine the frequentist objective probability. If instead all past and future cases were included, then narrowing in on the frequency in the reference class of days-like-today with only one member (namely, today) would yield a “frequency” (either 0, if in fact it doesn't rain tomorrow) or 1 (if it does) that is splendid for guiding credence about rain tomorrow. But nobody wants to salvage *PP* by making the concept of chance degenerate into that of truth/falsity.

<sup>13</sup>I have not discussed so-called “hypothetical frequentism” because it seems to me that such accounts usually amount to propensity theories, once they are fully spelled out. What a Humean wants is to identify chances with some actual facts – aspects or patterns, of some sort, in the huge panoply of actual events, able to play the chance role as specified in *PP*. If, *after* identifying the chances as something actual, one wishes to go on and assert that, in addition, they inform us about *what limiting frequencies would result if the antecedent conditions could be repeated infinitely*, that is one's own business. I personally don't see the need for this metaphysical extravagance.

There is no space here to go into the details of Lewis' account and the many ways in which (I believe) it goes astray. What I do wish to keep from his account is the idea of chances supervening on actual occurrences, and the idea of *systematic patterns* to be discerned in those occurrences, patterns that may be something more than just actual frequencies, and which can sensibly play the role of chance defined in PP. Just to give the simplest example of how this may work: The overall pattern of events may exhibit the kind of behavior patterns known as Newtonian mechanics (for middle-sized objects in certain circumstances). That fact, plus the symmetry of objects like coins and dice, gets us almost all the pattern-facts we need to see that the chance of heads on flipping a coin is 1/2 (and 1/6, respectively, for the die). The further fact we need is an aspect of the overall pattern of events that is truly crucial to the existence of objective chances. We might call it the "micro-stochasticity of events". In the case of coin flips, what this refers to is the fact that there is a nice random-looking distribution in the size, angle, etc. of the initial impulses given to coins in ordinary coin flips. If coin flipping is basically a Newtonian phenomenon, then it is the random-looking distribution of these initial impulses that makes coin flips display the approximate 50/50 distribution we rely on.<sup>14</sup>

The stochastic-lookingness of initial conditions, boundary conditions, influences from outside, etc., is such an important aspect of the overall Humean pattern of actual events that it deserves a title, and I propose to call it the Stochasticity Postulate. I call it a "postulate" because we don't *know*, for a guaranteed fact, that we can rely on it everywhere and at all times.<sup>15</sup> But it is as well-confirmed as anything in our scientific world-picture, and we rely on it to make many of our machines – nomological or otherwise – function predictably and reliably. It is not restricted to microphysics; for the purposes of economics, the car-buying decisions of consumers may supply the micro-stochasticity that is needed for an efficient model of new-car-delivery to work adequately well. Exaggerating only slightly, we might put the Stochasticity Postulate like this: all over the place, at all sorts of levels, events are nicely random-looking. It is this fact, above all else, that grounds the existence of Humean objective chances.

But the stochastic-lookingness of events does not, by itself, give us stable and reliable objective chances of the kind that could (ideally) serve to guide belief as per PP. We need in addition a stable structure or set of conditions that utilize this stochasticity, in constrained ways, to generate stable probabilities. As I said earlier, we need proper chance set-ups. Generally I will follow Cartwright (1999)'s terminology and describe these setups as probability-generating *nomological machines*. A nomological machine is a stable arrangement of things, with appropriate shielding as needed, that generates a regularity. A probability-generating nomological machine (or *stochastic nomological machine*, SNM, as I propose to call them) is a well-defined setup or arrangement of things that produces outcomes with a well-defined probability. It is thus something over and above mere superficial Humean "laws of association" (i.e., actual frequencies), and will therefore violate Pearson's admonition to avoid

---

<sup>14</sup>If coin flipping is not best thought of as a deterministic Newtonian process (e.g., if quantum interactions between coin and air molecules play an important role), then *other* sources of micro-stochasticity may be involved. But either way, it is the random-lookingness of influences at the micro-level (relative to the coin) that account for the actual statistical behaviours of coins.

<sup>15</sup>See the film "Rosenkrantz and Guildenstern are Dead" for a lovely example of the breakdown of this postulate.



entanglement with theory – but only, I think, to an extent that is both harmless and unavoidable.

The best examples of SNMs are, naturally, classical gambling devices, so let us look at a few of them to illustrate the main points.

*1. The coin flipper.* Not every flip of a coin is an instantiation of the SNM we implicitly assume is responsible for the fair 50/50 odds of getting heads or tails when we flip coins for certain purposes. Young children’s flips often turn the coin only one time; flips where the coin lands on a grooved floor not infrequently fail to yield either heads or tails; Persi Diaconis was alleged to be able to reliably achieve statistics far from 50/50 when flipping a coin in an apparently normal way (and he is no doubt not the first person to achieve this); and so on. Yet there is a wide range of circumstances that do instantiate the SNM of a fair coin flip, and we might characterize the machine roughly as follows:

- i. The coin is given a goodly upward impulse, so that it travels at least a foot upward and at least a foot downward before being caught or bouncing;
- ii. The coin rotates while in the air, at a decent rate and a goodly number of times;
- iii. The coin is a reasonable approximation to a perfect disc, with reasonably uniform density and uniform magnetic properties (if any);
- iv. The coin is either caught by someone not trying to achieve any particular outcome, or is allowed to bounce and come to rest on a fairly flat surface without interference
- v. If multiple flips are undertaken, the initial impulses should be distributed randomly over a decent range of values so that both the height achieved and the rate of spin do not cluster tightly around any particular value.

Two points about this SNM are worth mentioning right off. First, the characterization is obviously vague. This is not a defect. If you try to characterize what is an *automobile*, you will generate a description with similar vagueness at many points. This does not mean that there are no automobiles in reality. Second, the last clause refers to a “random distribution” in the initial impulses, and this might seem to be cheating, or creating some sort of vicious circularity. But in fact this is not the case. “Random” here simply means “random-enough looking” and has nothing to do with a mysterious “process-randomness” that fails to supervene on the actual happenings. For example, we might instantiate our SNM with a very tightly calibrated flipping machine that chooses (a) the size of the initial impulse, and (b) the distance and angle off-center of the impulse, by selecting the values from a pseudo-random number generating algorithm. In “the wild”, of course, the reliability of nicely randomly-distributed initial conditions for coin flips is an aspect of the Stochasticity Postulate.

*2. The biased coin flipper.* Here I will describe a proper machine, and not worry whether Persi Diaconis or other practitioners of legerdemain fit the description. Suppose we take the tightly-calibrated coin flipper (and “fair” coin) mentioned above, and: make sure that the coins land on a very flat and smooth, but very mushy surface (so that they never, or almost never, bounce); try various inputs for the initial impulses until we find one that regularly has the coin landing heads when started heads-up, as long as nothing disturbs the machine; and finally, shield the machine from outside disturbances. Such a machine can no doubt be built (probably has been built, I would guess), and with enough engineering sweat can be made to yield as close to chance = 1.0 of heads as we wish.

This is just as good an SNM as the ordinary coin flipper, if perhaps harder to achieve in practice. Both yield a regularity, namely a determinate objective probability

of the outcome heads. But it is interesting to note the differences in the kinds of “shielding” required in the two cases. In the first, what we need is shielding from conditions that bias the results (intentional or not). Conditions i, ii, iv and v are all, in part at least, shielding conditions. But in the biased coin flipper the shielding we need is of the more prosaic sort that many of our finely tuned and sensitive machines need: protection from bumps, wind, vibration, etc. Yet, unless we are aiming at a chance of heads of precisely 1.0, we cannot shield out these micro-stochastic influences completely! This machine makes use of the micro-stochasticity of events, but a more delicate and refined use. We can confidently predict that the machine would be harder to make and keep stable, than an ordinary 50/50 -generating machine. There would be a tendency of the frequencies to slide towards 1.0 (if the shielding works too well), or back toward 0.5 (if it lets in too much from outside).

3. *The radium atom decay.* Nothing much needs to be said here, as current scientific theory says that this is a SNM with no moving parts and no need of shielding. In this respect it is an unusual SNM, and some will wish for some explanation of the reliability of the machine. Whether we can have one or not remains to be seen.

In each of these cases we are able to describe a repeatable set of conditions that constitute the chance setup or SNM, and give at least some reasons for expecting it to yield a fairly reliable regularity. Sometimes the reasons may be expressed in causal terms; I think Cartwright expects this to be the case most, if not all the time. The reasons may also be grounded partly or wholly in what we take to be laws of nature, as is the case in the biased flipper (presumably modellable decently well with classical mechanics) and the radium atom (where the decay half-life follows from laws of quantum mechanics). This may seem to undermine the Humean credentials of objective probabilities. But there are two responses to this worry. First, there is an ineradicable link (or constraint) between the chances and the actual outcomes, at least when the numbers are high enough: had 99% of all coin flips in history landed heads despite the apparent satisfaction (in a huge variety of different ways) of conditions i - v, we would have to say that the objective chance of heads is 0.99, not 0.5.<sup>16</sup> The objective chances may be different from the actual frequencies, to some extent, in light of features of the chance setup (such as physical symmetry, presumed random-looking distribution of initial and boundary conditions, and so on), but not *greatly* different, at least not when the numbers are high. This constraint arises automatically from the need to satisfy PP. If chance is to be a good guide to belief, and 0.99 of all coin flips in world history land heads, then the chance had better be 0.99 too, or very close to it.

Second, while we may need to use causal and/or law-talk in describing our reasons for believing in the reliability of an SNM, we are not committed to any *particular* metaphysical account of these notions. Lewis, for example, offers accounts both of laws and of causation that satisfy his view of Humean supervenience. While I do not subscribe to those accounts, the point remains that this Humean account of objective probabilities leaves it an open question what account of causation or of laws is best (if any is needed at all, in the end). The notion of a SNM does not come loaded with any particular anti-Humean notion of probabilistic propensity. Indeed, in most of the cases I can think of, causal talk covers mainly the “deterministic” part of the

---

<sup>16</sup>In Hofer (1997) I argue that any Humean approach to chance is obliged to take this stance, denying the possibility of radically improbable outcomes for large sets of chance events such that the actual frequencies diverge strongly from the alleged objective chances.

workings of an SNM (e.g., for coin flips: what goes up, comes down because of gravity), while the part of the description that justifies the stochasticity, and the expectation of a stable probability, adverts mainly to the “randomness” of the inputs to the SNM from outside (force of the impulse on the coin or roulette ball, disturbing effects from random wind forces, etc.). And that is all just part of the Humean-acceptable pattern of actual events.

*4. Inflation > 6% in the UK economy.* What is the probability that inflation will exceed 6%, next year, in the UK? This example, as well as 1. and 2. above, is discussed in Cartwright (1999).<sup>17</sup> But this is an example of something that is *not* a proper chance set up, not a SNM. Why? It simply has none of the elements of such: no repeatable structure that it is reasonable to expect to generate a stable probability. If we are to correctly ascribe an objective probability, it would have to be based on a stable, enduring structure that is such as to reliably yield that probability. But over the years, both the meaning and structure of these notions (UK, inflation) changes greatly. There is no reason to think that any SNM is out there, waiting to be discerned by economists, that in fact grounds an objective  $P(I > 6\% | UK)$ .<sup>18</sup> If the UK economy lasts a few hundred years more and if we *could* see the statistics for all years, my guess is that there would probably not be any stable regularity discernible in them (e.g., inflation > 6% in approximately 3 out of every 18 years). Certainly, we cannot discern any reason why there *should* be such a regularity. There is no chance-generating nomological machine here, and so there is no objective chance.

Unfortunately, the vast majority of statistics that we can gather in economics, medicine, and other sciences – even statistics that we feel are important, and that we wish to understand and control – will be like this example, and not like the first three. There are many more statistics in the world than objective chances. For statistics can be seen everywhere, but genuine nomological machines – stochastic or otherwise – are much more rare.

Now we can return to the topic of causal laws and effective strategies.

**3. Re-thinking the examples.** How does spraying oil on swamps prevent malaria? We know the answer very well:<sup>19</sup> Particles of oil kill mosquitos when ingested (or when they land on larvae, perhaps); mosquitos are the carriers of the malaria virus; when the swamps are sprayed, some mosquitos should be killed (or larvae killed); so there should be fewer mosquitos around afterward; hence fewer mosquito bites; hence fewer bites by malaria-carrying mosquitos; hence fewer cases of malaria. Each of these steps makes common-sense causal sense; but each is also merely probabilistic, in some sense. The oil may kill more or fewer mosquitos, but is unlikely to kill all; fewer mosquitos should mean fewer bites, though of course that depends on how active the remaining mosquitos are; fewer bites should mean fewer bites transmitting malaria,

---

<sup>17</sup>Cartwright (1999) borrows the first two examples from a discussion by Mary Morgan and David Hendry (1995).

<sup>18</sup>This does not mean that there are no SNMs in economics generally. And with some work, we can imagine a fictional setup for the UK economy that might constitute a genuine SNM for inflation of a certain level. But the actual world is not such a setup.

<sup>19</sup> Actually, I am making this up, I do not *know* how this process works. But I assume some people do. More importantly, the true story will have a number of stages to it, like my possibly-fictional reconstruction here.

though again it depends on precisely how active the malaria-carrying mosquitos are, and how frequently their bites do in fact transmit the virus; and so on. In fact, at any of these stages if things don't *happen* to go the way one expects (due to chance, or unusual initial conditions if one prefers to think of it deterministically), then the oil-spraying may *fail* to reduce the rate of malarial infection. The problem for Cartwright's (1979) picture is this: this mooted failure of the "right" statistical relation to obtain is not due to any "missing" causal factors for malaria that we have failed to hold fixed. "Bad luck with which mosquitos survived" and "Bad luck with which mosquitos bit more" are not causes that Cartwright can recognize, or declare to be part of "the causal laws of the universe".

This sort of bad luck might have turned out to be universal in the whole reference class of oil-sprayed-swamps. More realistically, it might just happen in a few cases, leading to (say) the probability of malaria going up in one or more of the reference classes  $\wedge K_i$  mentioned in *CC*. Let's suppose this happens for the class in which we hold fixed {Don't drink quinine, use bug repellent, European ancestry, etc.} Then contextual unanimity (Dupre's term for the demand that the probability change in the same direction in *all* reference classes  $\wedge K_i$ ) would fail, contrary to what Cartwright thinks is possible in (1979).<sup>20</sup>

But we need to dig deeper into several aspects of this case, especially the probabilities. For there would be an obvious response to this example, if Cartwright were supposing the probabilities in *CC* to be the "true" probabilities, identical to the "real" propensities of systems of such-and-so type. The response would be: well, these statistics just don't count; they don't reflect the *true* probabilities. *CC* is still true, but only of the true probabilities.

But as we noted earlier, Cartwright does not hold with such things (which are in effect chances-as-metaphysical-propensities), and she is right not to.<sup>21</sup> Instead, as we saw, for Cartwright in 1979/1983, probabilities are just actual frequencies meeting certain tests. And that's not bad, from my perspective: it is better than invoking mythical propensities or hypothetical frequencies, and some such actual frequencies are indeed objective probabilities. But not all, by any means! And surely the frequencies of malaria infection, *in the tiny populations in which all these K-factors are held fixed*, are not – or are not all – genuine Humean OC probabilities. In fact, they are unlikely even to meet Cartwright's criteria of stability and so on. Nor, I would guess, will they in general meet the criteria for a chance-generating SNM.<sup>22</sup>

One might think that, once all the causal factors are held fixed and the situations of the classes  $\wedge K_i$  clearly defined, then a stable SNM must surely be the result.

<sup>20</sup>In Dupré and Cartwright (1988) and Cartwright (1989) she does allow that failures of contextual unanimity may occur, and she does not endorse it except where it reflects the presence of a stable causal capacity. But the reason for its failure in these works is "mixed" causal powers on the part of some causes, or interaction, rather than statistical bad luck.

<sup>21</sup>Some philosophers think of the "true probabilities" not as metaphysical propensities, but rather as parts of scientific models of certain situations. This is not the place to discuss the virtues of such a proposal, and how it may differ from the Humean account I favor; what matters here is that we are looking at situations for which we have *no* model, nor any reason to think that (in a non-trivial sense) we can have one.

<sup>22</sup>In (1989) probabilities are no longer actual frequencies for Cartwright, and instead are something more idealized. She does not give an overt account of what they are, but the perspective of *Nature's Capacities* may be seen as moving toward the view adopted in *Dappled World*.

Unfortunately, this is just not so, in general. The patterns among events at the macro- and micro-levels, relative to mosquito bites and so on, may not display any systematic regularity that entails, e.g., that in a particular homogeneous reference class, the chance of infection should be .046 rather than .054. The chance is supposed to reflect what *would* occur were the reference classes sufficiently large; but the actual patterns of events simply are not enough to dictate an answer to this question. This is in contrast with systems such as gambling devices, where the physical symmetries and the even distribution of initial and boundary conditions found all over the place in nature do dictate well-defined chances.

What if there really are no probabilities out there, in the reference classes holding fixed all the causally relevant factors? Then the project of deriving causal laws from probabilistic data is impossible to even begin. That is no big loss, however, since we already knew that we had to know, ahead of time, what all the other causally relevant factors are (to hold them fixed), in order to prove that a given factor is indeed a cause. Our knowledge of “the causal laws” has to be anyway, on Cartwright’s (1979) view, before we can look to probabilistic data to help complete it. As Cartwright has often stressed, the situation may not be so bad: if we feel we know enough about the causal structure(s) at issue, we may be able to use randomized controlled trials to discover whether C raises the probability of E in various subpopulations that we can’t examine individually. But this only salvages the utility of the CC-based method on the assumption that the method is applicable in the first place, i.e., assuming that *all the statistics we need to look at correspond to genuine objective chances*.<sup>23</sup> Unfortunately, this will not be true in general, at least if we understand objective chances in the way that I advocate here, or that Cartwright advocates in *Dappled World*.

What is it, for there to be a causal law of our universe that C causes E? CC cannot now be taken as part of the answer to this question. Contextual unanimity may fail (in actual statistics) for non-causal reasons, as well as the causal reasons recognized by writers on causation; and, much more importantly, the objective probabilities invoked in CC may simply fail to exist.<sup>24</sup> Instead, we must fall back on the answer from *Nature’s Capacities and Their Measurement* (1989): it is a causal law that C causes E just in case some *c*’s do, on some occasions, by virtue of being *c*’s, cause *e*’s.<sup>25</sup> And this singular causation concept is one that notoriously resists all attempts at further definition (or analysis), though I think we can say two things about Cartwright’s views

---

<sup>23</sup> By “correspond” here I just mean that two conditions are fulfilled: (a) the objective chances do, in fact, exist; and (b) the statistics being looked at are appropriately close to them. Typically causal searchers hoping to infer causal relationships from statistical data only consider condition (b), and deal with it by making it an unabashed, optimistic starting assumption.

<sup>24</sup> Of course, not every theorist of probabilistic causation defends the kind of contextual unanimity found in CC, and there are alternatives to CC that weaken the requirement. They do so, however, to handle cases like the “mixed causal capacity” of birth control pills to both cause and prevent thrombosis. That is not the sort of problem we are looking at here. The problems arising from either non-existence of objective probabilities, or (if one takes the probabilities to be by definition the actual statistics) accidentally misleading statistics, affect these other versions of probabilistic causation just as much as they do the views of Cartwright (1983).

<sup>25</sup> “A generic claim, such as ‘Aspirins relieve headaches’, is best seen as a modalized singular claim: ‘An aspirin can relieve a headache’; and the surest sign that an aspirin can do so is that sometimes one does so.” (1989), p. 95.

on the matter, based on her later writings. First, at least in some cases the common-sense counterfactual is true: “If  $c$  had not occurred,  $e$  would not have occurred.” Second, again at least sometimes,  $C$  will be part of an INUS condition for  $E$ , hence a particular  $c$  may occur along with the rest of a set of circumstances instantiating an INUS condition, and jointly necessitate  $e$ . But since we know we can’t really count on these explications being right all the time, basically we are down to this: that  $c$  singularly causes  $e$  is a primitive relation that we know how to recognize *sometimes*, and thank goodness, for without it we could never get science started.

While this last part is hard to deny – it is the core of Cartwright’s view of how science works, and it is the most true-to-life account anyone has yet offered – still it is possible to chip at it around the edges, and I think it is important to do so. Because again we seem to be retreating to a black-box perspective on causation, and again there are at least some cases where we know a lot about the internal mechanism. We could look back at the malaria case now, since we do know a lot about the internal mechanism. But it is an awkward example to use (and this in fact undercuts Cartwright’s story, to some extent), because it is hard to take seriously the singular-causation version of the mooted causal law. “Some oil-sprays do sometimes prevent malaria infections.” No doubt they do, but it is hard to read this in a singular-causation (prevention) way: *whose* malaria infections were prevented? Or at least, how many? Nor is an INUS condition reading very easy to put on the case, given that, as we saw, the effect is in some sense *likely*, but by no means *necessitated* to happen.

So instead let’s take Cartwright’s favourite example: some aspirin-takings do relieve headaches. Here too, we know a lot now about the mechanisms inside the black box causal statement.<sup>26</sup> Aspirin molecules float around in the stomach and get absorbed into the bloodstream. There, they mix thoroughly into the blood, and so some get pumped toward the brain. Because the molecules are small enough, they pass through the blood-brain barrier. The molecules then interact with the swollen vein and artery walls in the head, causing (by a yet-more-microscopic mechanism, which we will skip over) reduction in the swelling. The reduced swelling relieves the pressure that causes the pain.

That’s an awful lot of structure, hidden underneath a black-box-style singular causation statement. In fact it can be considered quite analogous to the oil-swamp-malaria case. At any of several stages of the story, the process relies on what are essentially statistical regularities (*not* brute cause-effect relationships, *not* necessitations supported by the (non-probabilistic) laws of nature): how many aspirin particles and of what size pass into the blood; how many pass into the relevant area of the brain; how many of these get into interactions that help reduce swelling; and so on. At each of these stages there is presumably a wide numerical distribution of the relevant events that may result, even when things go “normally”. And as in the malaria case, only perhaps more plausibly here, sometimes not enough reduction in swelling will result to cause headache-relief. And this will happen by mere chance, we may say, or by “hap”, or “just as a matter of random bad luck”.

When this occurs for the reasons just posited, we may advert to a useful metaphor and say that the cause “failed to fire” as a purely chance matter of fact. This can be misleading, though, in two respects. First, the metaphor calls to mind the (apparently) irreducible failure-to-decay that may be demonstrated, in a given stretch of

---

<sup>26</sup> As before, I am making this up, and as before the details do not matter for the philosophical points being illustrated.

time, by a radioactive atom. That is not a good comparison, since here we can in principle understand the failure to fire. The aspirin does what it always does, it is just that the micro-movements of its particles after swallowing happen not to be good enough to relieve the headache. We have a lot to say about what may have occurred, and none of it is black-box or irreducible.<sup>27</sup> Second, there is a temptation to assimilate this failure to fire to genuine “probabilistic causation”, thereby implying that there is some objective probability for this failure to occur (in a given population) at all times. But for the kinds of reasons already discussed above, there may in fact be no such objective probability.

The other way of thinking of the aspirin’s failure, consonant with the typical discussions of mixed capacities and interactions, would be to suppose that some *cause* prevents the aspirin from curing the headache. This may be wrong-headed as well. The aspirin may well not, on such an occasion, have been “prevented” from relieving the headache by any well-characterized factor whose causal power goes in the opposite direction, so to speak. Maybe it sometimes is, maybe even *most times* when it fails, it is. But it need not always be viewable that way. That is what my description was meant to highlight: the aspirin doesn’t necessarily fail because some more-powerful-headache-causer or aspirin-action-preventer wrestles it to the ground, but rather because at the micro-level, things just don’t happen to go the way they normally do.

When we look inside the black boxes of probabilistic causation, at least sometimes – and perhaps *every* time – we find a lot of stuff going on that is best described as a sequence of “causal” steps that rely on statistical regularities. Like the coin-flipping SNM, we can think of them as based around (fairly-)reliable statistical regularities that are treated either as unexplained, or as arising from the result of initial and boundary conditions given underlying natural laws. We may like to say that taking aspirin is an effective strategy for getting rid of a headache, because of aspirin’s “causal capacity” to relieve headaches. But underneath the metaphors of powers struggling and capacities firing, what’s really going on is the existence of some regularities that are stable and repeatable (-enough), which we exploit cleverly for our own ends.

Now let’s turn to effective strategies. In the malaria case, or the aspirin-taking case, I have been arguing that the CC-based causal law story breaks down upon close examination. *A fortiori*, it would seem, we can’t claim these are effective strategies on the basis of the truth of some causal laws. But that does not mean that these are *not* effective strategies! They probably are, in many or most circumstances. But in explicating why they are, we should avoid both talk of causal laws, *and* talk of specific objective chances at work in the strategies (either at the gross, desired-outcome level, or at the level of the underlying steps in the mechanism). Oil spraying and aspirin-taking are effective strategies not because there is an SNM (or NM) to be discerned in their working, but rather (merely) a *mechanism* that can be expected to work at least sometimes – perhaps often, if we are lucky.

Here, then, we see one big difference between NM’s and mechanisms in general: a mechanism need not give rise to a stable regularity. It simply has the potential, by virtue of its structure, to give rise to a certain outcome (or output) – when things go right. How often and how reliably they do go right is a separate question. Aspirin-

---

<sup>27</sup> Until, perhaps, we get down to the micro-chemical molecular interactions, which are in some sense quantum-mechanical – what matters here is not whether irreducible causation enters the picture somewhere deep down, but rather whether it is present at the level we start with.

taking is an effective strategy for curing a headache just because there exists the mechanism described above, that *can* (and does) function some of the time.

Before we say more about this view of effective strategies, I want to finish criticizing the causal law-based picture by looking at one more of Cartwright's central cases. Why isn't joining TIAA-CREF an effective strategy for extending your life? Well, actually, it might be, as Cartwright herself notes; and only a little imagination is needed to work out reasons why it could be. But let's suppose that on the whole it is not, in fact. Cartwright's (1979) story about why it is not goes like this: There is in the overall population a correlation between belonging to TIAA and having longer-than-average life. But once we partition the population into sub-classes in which we hold fixed the true causes of longevity (exercise? wealth? happiness? good genes? good diet? . . .) the correlation disappears. And at the level of singular causation, we can note: joining TIAA just never does cause longer life, in any individual case.

It is now clear what's deeply wrong about this story. First, the list of things that might be thought to affect longevity is too big, open-ended, and ill-defined for *CC* (or its strategies-directed correlate, from section 2 of (1983)) to be useful. And contrary to the "singular causes first" view of *Nature's Capacities and Their Measurement*, I would argue that there is no fact of the matter about whether, for example, 1 hour of hard exercise in the hot sun increases, decreases, or fails to affect my longevity. And the same could be said for a myriad of factors that may, *statistically*, be positively or negatively associated with lifespan.<sup>28</sup> But even at a non-singular level, the same problem arises: does exercising in the hot sun regularly or eating yoghurt daily cause greater longevity? There are reasons for answering yes, others for answering no, and still others for saying that there's no fact of the matter. (No SNM.) If we did manage to agree on a list of causes, and we partitioned the whole population up by homogeneity in these causes, our subpopulations would be too small to support genuine probabilities, on either my or Cartwright's early account of these. So the *CC* story just fails to make sense here.

Returning to TIAA at the level of singular causation: as in the malaria case, only much more so here, the mooted cause is so far removed from its effect, that (a) the notion of singular causing hardly seems decently applicable, but (b) if it is, then it is highly implausible that for *no-one* does joining TIAA actually increase their life expectancy. We can think of myriads of causal-counterfactual chains leading from joining TIAA to changes in lifestyle that *are* causes of increased life expectancy (to the extent anything is), and can imagine a person instantiating such changes. (Imagine a hard-drinking, smoking grad student who joins TIAA on getting her first academic job. She is sent a folio of information about how TIAA can help you get healthy by paying for your nicotine patches, subsidizing your health club membership, etc. . . .). But if the singular-causing story holds even once, then contrary to our starting assumption, it *is* a "causal law" that joining TIAA increases life expectancy – given the reading of *Nature's Capacities and Their Measurement*.

Finally, we should see how the pieces may fit together to offer a different, arguably Humean, picture of effective strategies.

#### **4. Humean effective strategies.**

---

<sup>28</sup>See Glennan (1997) for extended criticism of the singular-causation perspective on capacities found in *Nature's Capacities and Their Measurement*.



Cartwright's early way of talking about effective strategies may work well in a lot of cases, but in many others it tends to fall apart, as we have seen. The remedy, it seems to me, is to be even more stringent than *CC* in thinking about effective strategies, but stringent in somewhat different directions. Instead of looking for the complete sets of causal factors for a given effect, what we need to do is look for *mechanisms* or *nomological machines* – probabilistic or deterministic – that “produce” the effect. Where we can create, or discern in nature, a mechanism or a NM for a given effect, there we have a strategy for bringing it about. Where we can't find one, there we don't have an effective strategy, at least not one we have reason to think we can rely on.<sup>29</sup> We may have statistical regularities, and we may follow our temptation to base an effective strategy on the regularities. It may even work successfully in some cases. But that is just getting lucky; without a mechanism or NM, we are shooting in the dark.<sup>30</sup>

By contrast, if you have a mechanism or NM for producing an effect, you don't need to know *all* the causes and preventatives of the given effect. Instead, the mechanism/NM builds in “shielding” from interference, of two kinds. First, overt shielding from known disturbances, about which I have nothing in particular to say. Second, shielding by random initial and boundary conditions: the NM relies on nature's own fortuitous tendency to distribute uniformly the microscopic factors that might skew the results in undesired ways. This is of course analogous to the way in which human experimenters try to control for unknown skewing factors by randomized controlled experiments. But at a relatively microscopic level of description, Nature usually takes care of the randomizing for us, and that – part of the Humean supervenience pattern in the actual events – is a key fact around which many of our mechanisms and NMs are based.

A good example of such an NM to illustrate the role of nature's randomizing is the classical statistical-mechanical model of something like an ice cube being used to cool down a tepid drink. The model may not correspond to reality – it doesn't have to, to serve its illustrative purposes. But it may well so correspond, in its salient features.

What could be a more effective strategy for cooling down a tepid drink, than dropping a couple of ice cubes in it? Few things in this world are so reliable. But according to the classical stat-mech model, the strategy works not because of iron deterministic law, nor because of primitive causal powers of ice cubes to cool. Instead, the micro-motions of the liquid and the ice cubes are going to be *almost always* such

---

<sup>29</sup>Here I am joining NM's and mechanisms, which are closely related things but not the same. The difference: a mechanism brings about a result fairly reliably, but the result need not be a *regularity*, whether statistical or not. A NM generates a regularity (fairly reliably).

<sup>30</sup>It might be thought that this is too strong, and that surely if we have a positive statistical relationship between C and E (perhaps holding fixed some possibly-relevant and easy-to-measure further variables), then we have *prima facie* evidence that increasing C is an effective strategy for producing more E. I would deny even this *prima facie* claim. There are myriads of positive statistical relationships out there in the raw data, even ones that obtain given the stipulated constraints. Few of these will we ever measure, but they are there, and few of *them* correspond to genuine effective strategies. If, in practical experience, the kinds of variables we *do* measure and test in these ways turn out often to reflect causal connections, that is because we had good reason to suspect, prior to doing the statistical tests, that such a relationship might obtain. And such suspicions most often come from common sense and antecedent causal/mechanical knowledge (i.e, from suspecting there is the right sort of NM or mechanism to be found), not from noticing a statistical correlation.

that the future evolution of the system (ignoring outside influences) involves approach to equilibrium, with the equilibrium temperature being of course cooler than the initial temp of the liquid. This is the story, ignoring outside influences (and many other complications). But we should not ignore the environment: for this to be a good nomological machine for cooling drinks, it must be adequately shielded from outside heating. It must also be shielded from *coincidentally unfortunate boundary conditions (BCs)*, i.e., bumps from the outside that just happen, by bad luck, to be such as to keep the liquid + ice mixture moving *away* from equilibrium rather than toward it. But *we* don't provide this second kind of shielding; nature does that for us, via the reliable *typicality* and *randomness* of ICs and BCs to be found at the (relative) micro-level. Like any NM, it may on some occasion fail, but this one is a pretty good one compared to most that we devise. And notice one key point: the randomness (random-lookingness) of the micro-movements of molecules that is a key aspect of the pattern of actual events for a Humean account of objective chance is also the crucial to the functioning of this NM.

Laws and initial conditions underlie this SNM, not causes or capacities. Of course, this model of the situation relies on an ontological picture (billiard-ball style molecules interacting by action-at-a-distance forces, under Newtonian mechanical laws) that Cartwright would find incredible. And it may indeed be nothing more than a fiction. But if it is, it is a fiction that still works remarkably well at modelling one of nature's most reliable regularities. In light of it, and other examples that we could multiply indefinitely (the coin-flipping machine being another, for example), the claim that we should resign ourselves to causal fundamentalism in understanding our NMs seems premature.

Let me illustrate the NM-based view of effective strategies with a final case, the infamous heart-disease and exercising example, to point up how it is true to what we actually do when looking for real mechanisms in nature. The initially observed correlation between smoking and having less heart disease does not prompt us to immediately seize on smoking as an effective strategy for reducing heart disease (though it might be, if somehow smoking induces people to exercise who otherwise wouldn't). Rather it induces us to look to see if there might be a NM or mechanism linking smoking to reduced heart disease. There are two sides to this task. First, we may conduct further statistical studies to try to verify whether there really is such a mechanism at work, tests that give evidence *that* such a linkage exists without doing much to reveal *what* the mechanism is. For such tests the danger of misleading correlations is always there, and the implicit advice of *CC* – hold fixed known causally relevant factors as much as you can, and randomize – is of course correct as far as it goes. Second, we may directly test possible NM mechanisms *via* the hypothetico-deductive and other methods. We can no doubt immediately think of several ideas to test out: for example, nicotine might enter the bloodstream and have the effect of dissolving small clots inside the arteries, making the blood run more freely. Testing this might be more or less tricky, and depending on how it was done, might have more or less risk of deception via the exercise-heart disease link (or other correlations). But some tests of potential NMs might be fairly easy to do, and not have to rely on inferences made from mechanism-blind statistical studies.

The search for a NM or mechanism linking smoking to heart disease (or its prevention) is not a search for a causal law (as implicitly defined by *CC*), nor is it a search for singular causings. In these senses, I would say that the search is Humean-neutral: it does not imply causal fundamentalism at the level of the relationship under

study, nor does it imply that a non-Humean notion of causation is *not* needed, at the lower level where we describe the workings of mechanisms or NMs.

**5. Conclusion.** What I hope to have shown in this rambling discussion is that the framework of Cartwright's early discussions of causal laws and effective strategies is in many ways fragile, and that a different view of effective strategies is possible that makes use of her later concept of a nomological machine (and/or Glennan's concept of a mechanism). This view fits nicely with the Humean approach to objective probabilities that I advocate, which is agnostic about causation. The alternate view of effective strategies based on NMs is not meant to be Humean-sanitized, *vis a vis* causation: (i) I have not tried to revive the Humean project of defining "causal" facts purely in terms of statistical relations, a doomed project; (ii) as we dug into the various mechanisms by which causes such as aspirins effectively produce effects such as headache-relief, we had causal talk popping up frequently at the lower levels of description. But this doesn't mean that regularities (statistical, law-like or merely universal) are not enough to reconstruct what is going on, or that we need to fall back on some notion of causal capacity or causal law as a primitive, at the lower level. It means that the question is left open, and we can remain agnostic. My personal suspicion is that talk of causal capacities and causal laws can be replaced by Humean NM's all the way down to a level where all that is left are iron deterministic physical laws and fortunate accidental regularities.<sup>31</sup> Lots of fortunate regularities, which underlie at least as much of the predictability and stability of nature that we count on as the iron laws do (as we see in the ice-melting example). But this is not something I claim to have shown. Instead, what I hope my discussion has shown is that, *at the level of the original "causal laws"* that Cartwright wished us to accept, we can reject the need for any such things as primitives, and also reject their correlate singular-causings (taken, again, as primitives), and thereby make room, *at the level of these causal laws, relations and effective strategies*, for a more Humean approach to succeed.

Finally, let me stress that most of the points I have tried to make here are in harmony with Cartwright's most recent work on causation and probability, in *The Dappled World* (especially chapters 4, 5 and 7). The lessons I would wish to draw might be put this way: talk of causal laws should perhaps be avoided where possible, and the fact that causal capacities exist because of underlying mechanisms deserves more emphasis and investigation. Or more bluntly: it is better not to be too much of a causal fundamentalist.

A further result of these considerations seems to me worth mentioning (one that has been implicit in Cartwright's work from *How the Laws of Physics Lie* onward, but especially strongly in *The Dappled World*): the proposed methodology of trying to read off *useful* causal conclusions (hence effective strategies) from purely statistical data is *really hopeless*. In the first book, it proved hopeless because to decide that C was a cause of E (and hence a handle for increasing the level of E, at least in principle), you had to know *all the other causes of E first*. The methods of Spirtes, Glymour and Scheines (SGS, 1993) are meant to help one partially circumvent that problem, and they build in all sorts of idealistic features to their causal graph-systems to try to make it work (e.g. CMC, faithfulness). In *Nature's Capacities and Their Measurement* and *The Dappled World*, Cartwright argues very effectively that these assumptions are implausible, for the real

---

<sup>31</sup> Stuart Glennan seems to have a similar suspicion in his discussion of the mechanism-based view of causation. See (1996) section 4.

world in general. That is bad enough already. But perhaps the worst problem of all is one she doesn't sufficiently highlight: *most of the objective probabilities one needs as input simply don't exist*. There are SNMs in the world, including some we don't ourselves make. But they are hardly ubiquitous. And where they don't exist, the methods of Pearl and SGS may be literally inapplicable. Unfortunately, such methods are most likely to be needed and desired in precisely the sorts of fields (like macroeconomics) where it is extremely implausible that all the probabilities needed really exist. In those areas, what we have are at best what I would call *mere statistics*, not probabilities.

The difference is crucial. When you have a set of variables that are all connected by NM-like stable structures (or, using the SGS terminology, whose values are *generated by a causal graph*), there is at least some *prima facie* plausibility to the claim that the data will conform to the causal Markov condition and to faithfulness. But for the messy domains of mere statistics, what sort of arguments can be given for these conditions? You can only argue for their holding *after* you know that the statistics were generated by a real causal structure (i.e., a set of NMs and/or SNMs). But the SGS methods are supposed to start with mere statistical data, and *search out* a causal structure hidden underneath. Evidently, this can only be justified if one assumes that *all* sets of statistical data we may get hold of, arise out of some causal structure or other involving just those variables. This is unlikely not only because we will often latch onto irrelevant variables (and leave out relevant ones), but also for the reason stressed by Nancy Cartwright: much of what happens may occur just "by hap". What this means for our purposes here is: much of what happens is not appropriately thought of as "arising from a causal structure among event-types" or "happening because of the causal laws of the universe."

The question the causal modellers need to address is this: can any argument for the potential utility of such methods be mounted, given that we must largely work with mere statistics that we *know* are not generally objective probabilities? Instead of further proofs of how we can get true conclusions from ideally perfect probabilities assuming very strong conditions such as CMC and faithfulness, what is needed is some exploration of this difficult question.<sup>32</sup>

---

<sup>32</sup> And, as Cartwright has stressed, there is another topic that might be more useful to address: the nomological machines that do exist out there in the world. "[A] causal structure arises from a nomological machine and holds only conditional on the proper running of the machine; and the methods for studying nomological machines are different from those we use to study the structures they give rise to. Unfortunately these methods do not yet have the kind of careful articulation and defence that Spirties, Glymour and Scheines and the Pearl group have developed for treating causal structures." (1999), pp. 134 – 5.

## Acknowledgments

I would like to thank Henrik Zinkernagel, Jordi Cat, Mauricio Suarez, José Díez, Paul Teller, and Stuart Glennan for helpful comments on earlier drafts of this paper. Special thanks go to Nancy Cartwright, whose extensive comments tried to set me straight about a number of issues and led to important improvements in the text. Needless to say, none of the above endorse anything written here.

## References

- Cartwright, N. *How the Laws of Physics Lie* (Oxford University Press, 1983).
- Cartwright, N. *Nature's Capacities and their Measurement* (Oxford University Press, 1989).
- Cartwright, N. *The Dappled World: A Study of the Boundaries of Science* (Cambridge University Press, 1999).
- Dupré, J. & Cartwright, N. "Probability and Causality: Why Hume and Indeterminism Don't Mix", *Noûs* **22**, 521 - 36.
- Glennan, S. (1996) "Mechanisms and the Nature of Causation", *Erkenntnis* **44**, 49 - 71.
- Glennan, S. (1997) "Capacities, Universality, and Singularity", *Philosophy of Science* **64**, 605 - 626.
- Glennan, S. (2002) "Contextual Unanimity and the Units of Selection Problem", *Philosophy of Science* ??, 2002.
- Hofer, C. (1997) "On Lewis' Objective Chance: *Humean Supervenience Debugged*", *Mind* **106** no. 422.
- Hofer, C. (2003) "The Third Way on Objective Chance", unpublished manuscript.
- Norton, J. (2003) "Causation as FolkScience", <http://www.philosophersimprint.org/003004/>
- Otte, R. (1985) "Probabilistic Causality and Simpson's Paradox", *Philosophy of Science* **52**, 110 - 125.
- Spirtes, P., Glymour, C., and Scheines, R. *Causation, Prediction and Search* (Springer-Verlag, 1993).