RESILIENCY, PROPENSITIES, AND CAUSAL NECESSITY *

CONSIDER laws of the form: "If a physical system is in state χ then the probability that it has property ϕ is α"; or "The probability that a physical system has property ϕ given that it has property $\chi = \alpha$." These laws may be generated by statistical treatment of an underlying deterministic process (e.g., statistical mechanics) or may be, to the best of our knowledge, basic (e.g., quantum mechanics). The physical system with property χ is said to have a *propensity to exhibit* ϕ *with probability* α. Propensities arising from the first sort of law (e.g., the propensity of a Brownian particle to migrate a certain distance in a certain time) are not essentially different from propensities generated by roulette wheels, etc.

Statistical laws and resultant statements of propensity cannot simply be saying that the limiting relative frequency of ϕ within the ensemble of systems that exhibit χ is α any more than a universal law "All *F*s are *G*s" can simply be saying that 100% of the *F*s are *G*s." There are the considerations of lawlikeness. We might get a relative frequency *by accident*. What is worse, we might not only get a relative frequency where there is no statistical law operative (a "spurious correlation") but we might get the *wrong* relative frequency even where a statistical law *is* operative. To make the point bluntly, suppose a statistical law sets $\alpha = 1/2$ and only one physical system ever exhibits χ and it is destroyed after one trial (for vividness, you might imagine a special roulette wheel built to oddball specifications). Then the limiting relative frequency is either 0 or 1. Less extreme examples of the same phenomenon emerge when we calculate *from* the law that there is a positive probability that in a finite number of trials the relative frequency will diverge from the probability.

To attempt to escape these embarrassments by conjuring up large ensembles is of a piece with trying to rescue lawlikeness for universally quantified material conditionals by appeal to "unrestricted scope"—and just as vain.

To avoid these unacceptable consequences, some writers have suggested that propensities be construed counterfactually, as hypothetical relative frequencies; or as dispositions to manifest those relative frequencies. Thus, Pr (ϕ given χ) $= \alpha$, if an infinite number

* To be presented in an APA symposium on Causation and Conditionals. Patrick Suppes, and Robert Stalnaker will comment; see this JOURNAL, this issue, 713/4 for Suppes' comment; Stalnaker's is not available at this time.

of physical systems were put into state x, the limiting relative frequency of ϕ among them would be α.

This view runs into some of the same difficulties as the foregoing, although in attenuated form. I believe that the probability of heads on a fair flip of a fair coin is $\frac{1}{2}$. Does this logically *entail* that if I were to flip a fair coin (by a fair flipping process) an infinite number of times I would not get one head after another? Again, I think not. Granted that the probability of "all heads" gets smaller as the number of trials increases and approaches 0 as the number of trials goes to infinity (hence the attenuation). *But we must not take zero probability as tantamount to impossibility in these contexts.* After all, the same assumptions that got us zero probability for "all heads" get us zero probability for *each* infinite sequence of heads and tails, and we cannot very well hold that they are *all* impossible! Therefore, appeal to laws of large numbers, and appeal to hypothetical situations to assure those large numbers, is not enough to guarantee a correct representation of propensities.

I want here to introduce a technical notion, the *resiliency* of a probability claim. Let a belief state be represented by a probability distribution $[\mathrm{Pr}_i]$ over a language, and let p be a sentence of the language. Then:

Def. Probabilistic Resiliency: The resiliency of $\mathrm{Pr}(p) = a$ in $[\mathrm{Pr}_i]$ is $1 -$ the maximum over j of $|\mathrm{Pr}_j(p) - a|$ where the $[\mathrm{Pr}_j]$s are the probability distributions got from $[\mathrm{Pr}_i]$ by conditionalizing on some sentence of the language consistent with p; $\sim p$.

Resiliency is a stability property, akin to the concept of robustness in statistics. A resilient probability is one that is relatively insensitive to perturbations in our belief structure. A treatment of conditional probabilities in the same spirit is possible:

Def. Resiliency for Conditional Probabilities: The resiliency of Pr (q given p) $= a$ in $[\mathrm{Pr}_i]$ is $1 -$ the maximum over j of Pr_j (q given p), where the $[\mathrm{Pr}_j]$s are the probability distributions got from $[\mathrm{Pr}_i]$ by conditionalization on some sentence that entails neither $p \supset q$ nor $p \supset \sim q$.

The concept of probabilistic resiliency is nicely illustrated by Richard Jeffrey's solution to Karl Popper's "paradox of ideal evidence."

Popper proposes the following problem. You are presented with a coin and are to assign rational degrees of belief to the probability that it will come up heads and the probability that it will come up tails. You are not sure that the coin is fair. You believe that there is some chance that it is biased, either toward heads or toward tails. But you have no more reason to think it is biased one way than the other. From "symmetry of ignorance," so to speak, you arrive at the conclusion that Pr(tails) $= 1/2$. Now compare this with the situation in which you toss the coin a great number of times and get about 50% tails; you examine the coin and find it physically symmetrical, etc. You now have a great deal of knowledge available that you did not have in the first case, and yet when asked your rational degree of belief that the coin will come up tails on the next toss, you will give the same answer: Pr(tails) $= 1/2$. The conclusion appears to be that your added knowledge is *simply not reflected* in your degrees of belief in the outcomes of the coin-tossing experiments. (Popper then wants to conclude that they can be reflected only in degrees of belief about *objective* probabilities: at this stage he is thinking about relative frequencies.)

In a slightly different context, Leonard Savage flirts with the idea of explaining the difference between subjective probabilities that we are "sure of" and those we are not sure of, by introducing second-order subjective probabilities:

> To approach the matter in a somewhat different way, there seems to be some probability relations about which we are relatively "sure" as compared with others. When our opinions, as reflected in real or envisaged action, are inconsistent we sacrifice the unsure opinions to the sure ones. The notion of "sure" and "unsure" introduced here is vague, and my complaint is precisely that neither the theory of personal probability as developed in this book, nor any other device known to me renders the notion less vague. There is some temptation to introduce probabilities of the second order.[1]

We appear to have a quite intuitive picture here of the situation in the Popper examples. In both situations the first-order probability of tails is $1/2$. But in the ignorance situation the *second-order* probabilities are spread out all over the spectrum for Pr (tails) $= x$ [though we may plausibly assume that the second-order probability weighed average for values of Pr (tails) $= x$; i.e., the second-order expectation, is $1/2$]. In the "ideal-evidence" situation, the second-order probabilities can be thought of as concentrated sharply at Pr (tails) $= 1/2$, so that Pr (Pr (tails) $= 1/2) = 1$ (or some close approximation to that situation). (Notice that *mathematically* the Savage picture may not be so different from what Popper has in mind. Savage thinks of both first- and second-order probabilities as subjective, while Popper thinks of second-order as subjective and first-order as objective.) Jeffrey points out, however, that we do not

[1] *The Foundations of Statistics* (New York: Wiley, 1954), pp. 57/8.

even need to ascend to second-order probabilities, in order to find the imprint of the additional evidence:

> I suggest at this point that your attention is being misdirected to A_n (the nth toss comes up heads) as a proposition to which your old and new belief functions must surely assign different values, if there be such a proposition. But as you both agree, prob (symmetric ignorance) and PROB (ideal evidence) shall both assign the value 1 here. Then this cannot be the locus of the difference. Nevertheless, there *is* a difference: prob and PROB will assign different values to any proposition $A(n)$ that asserts, concerning $n \geq 2$ distinct tosses, that all of them yield heads. To any such proposition PROB assigns the value $\frac{1}{2}^n$; but to the same proposition prob must assign a higher value, if you hope to learn from experience.[2]

In other words, the difference shows up in the first-order *conditional* probabilities. In the ignorance case, Pr (tails on toss 100) $= \frac{1}{2}$, but Pr (tails on toss 100 given heads on tosses 1 through 99) is nearer 0. In the ideal-evidence case, this conditional probability would stay at (or very near) $\frac{1}{2}$. Our "sureness" that Pr (tails) $= \frac{1}{2}$ is manifest as a reluctance to change Pr (tails) on various evidence, and this is mirrored by the constancy of value of Pr (tails on a given q) for various q. In a word, the ideal evidence has changed not the *probability* of tails on toss a, but rather the *resiliency* of the probability of tails on toss a.[3]

Now in the first "symmetric ignorance" case, our rational degree of belief in the proposition that the coin has a *propensity* to come up tails with probability $\frac{1}{2}$ is quite small, but in the "ideal-evidence" case, it is quite high. Let χ be a description of the kind of trial involved, and let ϕ be the property that it has a propensity to exhibit (e.g., tails). Then it seems plausible to take as our degree of confirmation that χ has a propensity to exhibit ϕ with Probability $= \alpha$, as the instantial resiliency with respect to an unexamined trial, a:

$$\text{Resiliency Pr } (\phi\, a \text{ given } \chi\, a) = \alpha$$

Likewise, a statistical law (or consequence of a statistical law) of the form: "The probability that a physical system has property ϕ given

[2] *Logic of Decision* (New York: McGraw-Hill, 1965), p. 184.

[3] These facts are not unrelated to the second-order approach. See discussions of the De Finetti representation theorem in Bruno de Finetti, "Foresight: Its Logical Laws, Its Subjective Sources," in H. E. Kyburg and H. E. Smokler, eds., *Studies in Subjective Probability* (New York: Wiley, 1963), in Savage, *op. cit.*, sec. 3.7., and in I. J. Good, *The Estimation of Probabilities*, Research Monograph No. 30 (Cambridge, Mass.: MIT Press, 1965).

that it is in state x is α" may be taken to be well confirmed to its degree of instantial resiliency, as specified above.

Resiliency over the whole language may be a requirement of unrealistic stringency. There is no unique answer as to which sublanguage resiliency must be evaluated over, for lawlikeness. Rather, we must again say that *the larger* the sublanguage over which we have high instantial resiliency, the more *lawlike* the statistical law. At one end of the scale we have statements like "the probability of death within a year given that one is an American male of age $65 = d$," which is extremely sensitive to auxiliary information, and whose resiliency is limited indeed. At the other end we have laws of radioactive decay, which have been tested under an enormous variety of circumstances and whose resiliency extends over a language of impressive scope.

We can, however, say something more about how broadly the language is taken over which resiliency is to be evaluated. It so happens that different choices of "scope" of resiliency correspond to different statistical properties which play an important role in the discussion of propensities.

Let us first compare the pre-quantum-mechanical status of the laws of statistical mechanics and their consequences (e.g., the random walk of a Brownian particle) with the current status of the laws of quantum mechanics and their results (e.g., laws of radioactive decay). In the former case we are thought to have a statistical situation only in virtue of our ignorance of the initial conditions of the system involved. Given the initial position and momentum of each particle involved, Newton's laws predict the evolution of the system deterministically. The quantum-mechanical laws on the other hand are thought to be basic (hidden-variable aficionados aside), and the statistical situation with respect to them is due not to ignorance of initial conditions, but to a "genuine metaphysical indeterminism." Some writers regard the propensities involved in the first sort of situation as bogus. They believe that in a deterministic universe there are no real propensities. Let us call this school the *indeterminists*. Pre-quantum-mechanics, the universe *was* thought to be deterministic. According to the indeterminists, then, in that knowledge situation we should *not* have thought of statistical mechanics as giving us *genuine propensities or genuine statistical laws*. Likewise, we should not have thought of roulette wheels, etc. as chance set-ups having genuine propensities. Other writers find the concepts of propensity and physical law equally applicable in both sorts of case.

Now, from the standpoint of the present treatment, the dispute can be seen as a dispute over how broadly to evaluate resiliency. The indeterminist is worried about the fact that, once we introduce descriptions of the microstate of the system into the language, the resiliency of the appropriate probability drops drastically. But those of us who are happy about talking about statistical laws and propensities with regard to classical statistical mechanics, are presumably impressed enough by resiliency over descriptions of the macrostate of the system. Similar remarks apply, *mutatis mutandis* to roulette wheels and other homely paradigm chance set-ups. At this point, it is not clear that *determinism* should be the point at issue here, rather than macro-description vs. micro-description or rough, everyday description vs. precise mechanical description (for the roulette wheel). But if indeterminism is *really* what the indeterminist wants, then it is clear that he should evaluate Resiliency of Pr ($\phi\ a$ given $\chi\ a$) $= \alpha$ over a language that includes resources for a complete description of the history of the world up to the time of occurrence of *a* (but not after that time). But this requirement is *very* strict, and its temporal asymmetry strikes me as a little odd in this context.

Another bone of contention among propensity theorists is whether the various trials produced by a chance set-up must be independent. Some wish to assume independence of trials, in order to use the strong law of large numbers in the afore-mentioned justification of limiting relative frequency. (The strong law of large numbers is that, if the trials are independent and identically distributed, the limiting probability that relative frequency of an outcome diverges from its probability is zero.) Independence is also a special case of resiliency; that is, resiliency of 1 over the results of other trials.

Finally, there is the question of whether the sequence of trials should be *random*. Randomness was introduced into the theory of physical probabilities by Richard von Mises. An infinite sequence (e.g., tosses of a coin) is to be *random with respect to its outcomes* (e.g., heads, tails) just in case the relative frequency of the outcomes remains unchanged in all subsequences got from the original sequence by "place selection." The intuitive idea behind place selection is:

> By place selection, we mean the selection of a partial sequence in such a way that we decide whether an element should or should not be included without making use of the attribute of the element.[4]

[4] *Probability, Statistics, and Truth* (New York: Macmillan, 1957), p. 25.

Von Mises tried to make this idea precise by identifying place selection as selection by a characteristic function: a function that takes as arguments initial segments of the sequence and as values 0 (signifying "next member not selected for the subsequence" and 1 (signifying "next member selected"). Of course, with the set-theoretic sense of function this will not do, for there are enough functions around to upset the claims to randomness of *all* sequences (excepting a few degenerate ones). Take an infinite sequence of heads and tails. Consider the function that maps an initial segment of the sequence onto 1 just in case the next element is Heads. There *is* such a function in the set-theoretic sense, although the way I have specified it may seem a little underhanded. And, provided there were an infinite number of heads in the original sequence, it will select out an infinite subsequence consisting entirely of heads. Likewise with tails. The problem is that the epistemic clause: 'without making use of the attribute of the element' has no restrictive role to play in this account.

What one can have in a nonvacuous way is a notion of randomness relativized to a certain class of place-selection functions. George Wald showed that, relative to an arbitrary denumerable class of place-selection functions, there is a continuum of random sequences. Alonzo Church suggests taking a particularly natural set of place-selection functions, the recursive ones.

It should be clear that these ideas of randomness are also closely connected with probabilistic resiliency. For every place-selection function there is a corresponding property that selects out the subsequence (e.g., the property of following the initial segment TTT or the initial segment TTTH or the initial segment TTTHHT or . . .) Resiliency of 1 over the instantiations of this class of properties will guarantee randomness relative to the associated class of place-selection functions. (And will coincide with it provided that the class of properties has the appropriate Boolean closure property.) The fact that so many key concepts are special cases of probabilistic resiliency should not be surprising. Resiliency of 1 comprises a very general and very fundamental concept of *invariance.* If resiliency of $\Pr(p)$ is 1 over a language, then $\Pr(p)$ is invariant over any situation consistent with p describable in the language (or, in other words, over all partitions of that region of logical space within which p is true, that can be generated by the language).

Looking at the other side of the coin, we might say that probabilistic resiliency is a natural generalization of von Mises' original definition of randomness. In fact, Hans Reichenbach objects that the direction suggested by Church construes the invariance too narrowly for the intended physical applications: "if a sequence possesses randomness of the von Mises-Church type there may still be *physical* selections that lead to a deviating frequency." [5] That is, some physical property (e.g., temperature below −200 degrees C) might select out a subsequence, or subensemble, which changed the frequency (and thus called for qualification of the associated physical law or propensity statement). Let us call Reichenbach's idea of invariance under selection of subsequence by an arbitrary physical property *physical randomness.*

The *absolute* concept of physical randomness is clearly in as much trouble as von Mises' original definition. Just what physical properties *exist* is a tricky physical question. If we take the extensional route of identifying physical properties with classes of physical events, we will have a bit of difficulty finding physically random sequences. But even without indulging in such dubious metaphysical identifications, we can see that the concept of absolute physical randomness is suspect. Consider the paradigm case of radioactive decay. We can select out subsequences with different relative frequencies simply by referring to the readings of detectors placed in the vicinity. There seems to me to be no reason to believe that we could not always, by referring to the results of physical measurements, select out subsequences with variant relative frequencies. The only sort of physical randomness that makes sense, then, is randomness *relative* to a given set of physical properties (in other words, randomness relative to a given language).

Of course, physical randomness as thought of by Reichenbach is a property of objective sequences of events in the world, and I have been presenting probabilistic resiliency as a property of our system of beliefs about the world. This is not to say that the account presented here is *anti*-objectivist. Probabilistic resiliency of 1 does not *require* belief in the existence of a objective, physically random, sequence of events. But it does not exclude the possibility that the resiliency is based in the belief *in* such a sequence. And, in fact, strong enough belief in physical randomness *will* guarantee the corresponding probabilistic resiliency.

The key tool in investigating these questions is the De Finetti representation theorem. A probability distribution is said to be *symmetric* for a sequence of trials just in case the probabilities are invariant over permutations of trials; equivalently, the trials are

[5] *Theory of Probability* (Berkeley: University of California Press, 1949), p. 150; I learned of this through Jose Alberto Coffa.

said to be *exchangeable.* (We have been assuming all along that "new" individuals—individuals not in our evidence base—are exchangeable.) De Finetti showed that the probabilities of an exchangeable series of events can be represented as a mixture (i.e., a weighted average) of probabilities of independent series of events. One way of using this mathematical result is to think of my subjective probabilities of events in such sequences as being weighted averages of objective probabilities; with the weights being our subjective probabilities that the corresponding objective probabilities are the true ones. Then my subjective resiliency of Pr $(Fa) = \alpha$ will approach 1 as my belief in the physical randomness of the corresponding sequence of events with Pr $(Fa) =$ Pr $(Fb) = \ldots = \alpha$ approaches 1.

The concept of resiliency has connections with a whole cluster of concepts associated with lawlikeness and causality. I can here only briefly indicate a few of these.

(1) *Shielding-off:* Suppose that e_1 and e_2, though not independent, became independent after conditionalizing on c, that is,

$$\Pr\,(e_2 \text{ given } e_1 \,\&\, c) = \Pr\,(e_2 \text{ given c})$$

Then c is said to *shield-off* e_1 from e_2. In a wide range of contexts it is plausible to assume that *shielding-off* holds the key to causal order: e.g., the current atmospheric conditions shield-off the dropping barometer from the impending rain.

To say that c *shields-off* e_2 from e_1 is to say that Pr (e_2 given c) has resiliency of 1 over e_1.

The plausibility of the connection between resiliency and causal ordering depends on the same sorts of assumptions that are required for the legitimacy of the second law of thermodynamics. But this reservation should increase rather than decrease the importance that resiliency has in this area.

(2) *Resiliency of Nonprobabilistic Statements and Rules of Acceptance:* The resiliency of a nonstatistical statement, p, can be identified with the limiting case of the resiliency of a statistical statement Pr $(p) = 1$. Then, the resiliency of p is the minimum of Pr (p given q_i) where the q_is are the propositions consistent with p in the language in question.

It is of some interest to see whether we can find probabilistic rules of acceptance that are *strongly consistent,* that is, which always lead to a consistent set of accepted sentences. The lottery paradox shows that high probability alone will not do. However, it is possible to show that:

Th. For any compact language, with a strictly coherent probability distribution on it, any rule of acceptance that requires resiliency $>.5$ will yield a consistent set of accepted statements.

(3) *Simple Nonstatistical Laws:* Some new light can be thrown on old issues concerning the confirmation of simple nonstatistical laws of the form "All *F*s are *G*s." Let us say that a numerical quantity *M* on laws of this form *supports prediction* iff Pr (*Ga* given *Fa*)[5] $\geq$ M (All *F*s are *G*s). A quantity that *supports prediction* is one that certifies only those laws which issue good inference tickets. Let us say that a quantity *M satisfies the equivalence condition* iff all logically equivalent laws[6] receive the same *M* values. Then, subject to certain restrictions, it can be shown that there is a unique quantity *M* which supports prediction and satisfies the equivalence condition, and that it is the resiliency of the material conditional which instantiates the law $R\,(Fa \supset Ga)$ over a suitable instantial language.

Thus two prima facie alternative positions:

I. That a law is expressible in terms of the material conditional, $\supset$, but that being well confirmed qua law involves a different *status* from merely high probability; and

II. That a law is *a bundle of conditionals which are not material* conditionals; e.g., a bundle of conditional probabilities,

can be seen to be compatible.

BRIAN SKYRMS

University of Illinois at Chicago Circle

SOME REMARKS ON THE CONCEPT OF RESILIENCY *

I AM sympathetic with the viewpoint Skyrms develops in his paper, but the formal characterization of the central concept of resiliency seems defective. For example, in a language having exactly two atoms, each with nonzero probability, the resiliency of the probability of each atom is 1. More generally, the resiliency

[5] For an unexamined *a*. I assume exchangeability for unexamined *a*s.

[6] Taking, for these purposes, "All *F*s are *G*s" as $(x)\ (Fx \supset Gx)$.

* Abstract of a paper to be presented in an APA symposium on Causation and Conditionals, December 30, 1977, commenting on Brian Skyrms, "Resiliency, Propensities, and Causal Necessity," this JOURNAL, this issue, 704–713.