How To Do Things With Causes[1]

NANCY CARTWRIGHT

UNIVERSITY OF CALIFORNIA–SAN DIEGO AND THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

Presidential Address delivered before the Eighty-Third Annual Pacific Division Meeting of The American Philosophical Association in Vancouver, British Columbia, on Friday, April 10, 2009.

1. An unnerving thesis

"It simply wouldn't be true to say, 'Nancy LD Cartwright...if you own a TIAA life insurance policy you'll live longer.' But it is a fact, nonetheless, that persons insured by TIAA do enjoy longer lifetimes, on the average, than persons insured by commercial insurance companies that serve the general public."[2]

This is from a letter urging me to buy TIAA life insurance: it's cheaper—college professors don't on the whole die young. It was the first sentence in my first paper on causality. Correlation, I remarked, is not causation. That is well known. But when I started my career you weren't allowed to believe it. Like a good Humean, you could believe in regularities and that's all. Causation—all the causation we were allowed—is correlation, just special correlations, ones with bells on.

This view had ramifications in decision theory. As a philosopher of science I cared about decision theory because I believed that science and philosophy were not only going to help us represent the world but to change it. That was, after all, the point of Selma, the march on Washington, the protest against the war, and the utopian communes. When it came to theories of how to decide what actions would be effective, the bar against causation had impact. Central versions of decision theory had been using the conditional probability of R on A—P(R/A)—to represent the probability that *if I were to do act A, R would result*. But P(R/A) is just correlation. And no bells will fix that. The probability that you will live longer if you buy TIAA insurance is high, but if you want an effective strategy for living longer, I urged in this early paper, you should do something not just associated with this result. You should do something connected by *causal law*.

This paper was one of the early forays in a growing movement to reinstate causality as a respectable member of our ontology. The paper was entitled "Causal Laws and Effective Strategies," and that marked out the thrust of the argument. Not all associations are equal. Some provide effective strategies, others do not; the difference is that the first are expressions of causal laws.

This kind of argument has been echoed repeatedly; one especially powerful voice for it now comes from the prominent invariance-under-manipulation camp of causality. I think this is unfortunate. For the first unnerving conclusion I argue for here is that causal laws are not, after all, much of a guide to action. If true this conclusion has far more than philosophical import. We—society as a whole and special institutions within it, from pharmaceuticals to the U.S. Department of Education—spend a vast amount of money, intellectual effort, and talent to establish causal laws, and we do so because we believe they are a guide to action.

2. What's in store

So that is the first thesis I will defend here:

Thesis 1: Causal laws are not an especially good guide to effective strategy.[3]

Although the problem is not identified in this stark form by others, there is a hugely popular fix on offer at the moment: *invariance*. This leads to my second thesis:

- Thesis 2: Invariance is no rescue.
  - Causal laws aren't especially invariant.
  - If invariance is available, we don't need causal laws.

I shall then unashamedly argue for an old idea of mine, the nomological machine, which I have not fussed over much for almost twenty years. But effective strategies matter and I believe these machines matter to effective strategies; we mustn't be fobbed off with just invariance or we'll never know how to do it.

Hence –

- Thesis 3: "Nomological machines" are the prime source of the special laws that are good for strategy.

And finally

- Thesis 4: "Nomological machines" are the real engines of change.

I take up again my old theme of nomological machines because I think we are doing it wrong—both in philosophy and in practice, and the two support each other. Evidence-based policy is now mandated at the local, the national, and the international level. What are we supposed to take as evidence that a policy will be effective? We are told: "A rigorous test that a causal law holds"; and when it comes to rigorous tests for causal laws, randomized controlled trials (RCTs) are all the rage. The evidence-based policy community knows though that RCTs are not enough because of the conventional trade-off between internal and external validity. RCTs are good on internal validity. A successful well-conducted RCT can establish with a high degree of certainty that the treatment causes the outcome in the population involved in the experiment under the conditions under which the treatment is administered in the experiment. The result has external validity if it holds for other populations under other circumstances. So practitioners as well as philosophers are going in for invariance. They are, among other things, investing heavily in devising statistical tests to check whether different sets of data support an invariance hypothesis.

I think this is the wrong way to go, or at least we should not be putting all our eggs in this basket. Testing for invariance is of limited use. We need instead to understand invariance, in particular to understand where it comes from, especially one major source, which I call the *nomological machine*. When we don't understand the machine and how it works, we can easily implement policies without ensuring the concomitant factors necessary to support the causal connection between action and output, and we don't get what we want. Worse, when we focus on invariance, looking to see when "the same old thing" will be true again, we ignore a far more powerful policy instrument, the possibility of building new nomological machines that give rise to new causal laws. So my worries about casual laws and effective strategies have real practical bite. Casual laws won't do the job and invariance is a weak fix. For effective strategies we need something more and different. I call this a nomological machine and have developed a theory about it. If you don't like the term or the theory, then let's get a better one. Because we can't do just with causal laws and invariances.

3. A nineteenth-century distinction

I am not alone in defending the importance of nomological machines for securing the kind of reliable laws employed in day-to-day strategizing. I read the defenders of *mechanisms*—my colleague Bill Bechtel[4] and Machamer, Darden and Craver (MDC)[5]—as doing the same. The problem is that "mechanism" is a term with a thousand meanings, which is why I made up a new term to pick out just what I thought we needed. And the special problem right now is that there is one other sense of "mechanism" that is sweeping the field and that, I maintain, cannot do the job: *mechanism* in the sense of "invariant causal law." Sometimes these two senses are conflated. So the story I am going to tell here is a tale of two senses of mechanism.[6]

I want to appeal to a nineteenth-century distinction to drive home my point, a distinction elaborated by the historian Norton Wise in a lovely paper on automata.[7] It has long been noticed, as Wise puts it, that nineteenth-century automata are typically women or "other sorts of uncanny or exotic creatures: talented children, blacks, acrobats, monkeys, magicians."[8] Automata perform mechanical actions that produce effects. The mechanical woman moves the pen and causes marks on the paper. But her action is merely mechanical. It is reliable. But the virtue of reliability is the flipside of a defect. The causal pattern is repetitive and unoriginal; little can be accomplished with it.

This, according to Wise, is a conventional nineteenth-century theme, which he illustrates across a number of discourses. Consider H.G. Wells's *The Time Machine*. The life of the Eloi is one of sameness and repetition. The Eloi is "a perfect mechanism" according to the time traveler. They are adapted "to a life of 'mere mechanical industry'." But they are not adapted to intelligent or creative action. The traveler remarks: "There is no intelligence where there is no change and no need to change."[9] Another case is Charles Babbage. Babbage divided objects of machinery into "1st, Those which are employed to produce power; and as, 2dly, Those which are intended merely to transmit force and execute work."

The nineteenth-century difference is between *engines* and *mechanisms*. Wise explains: "[E]ngine always implied productive power, as in referring to the differential calculus as 'the engine of analysis', while 'mechanism'...referred to a device for executing a particular form of motion, typically repetitive."[10] And, of course, the nineteenth century had it back to front with women as the automata (and workers on the factory floor) and men as the powerful and rational engines driving change and production.

The advocates of invariance—the view I worry about here—call their invariant laws "mechanisms." I shall allow them this term, but in Wise's nineteenth-century sense, of reliable, repetitive, unoriginal transmitters of work; the advocates of this view I call the "Mechanics." The items that really matter for understanding and for securing good strategies, I shall argue, are nomological machines including, if I am right in my interpretation of them, the mechanisms of Bechtel and MDC. These are the basic source both of stability and of change; following Babbage and Wise, I'll call these *engines*, and their advocates "Engineers." Before I begin the argument then, let me introduce the dominant players on the opposition side: James Woodward,[11] Judea Pearl,[12] Daniel Steel,[13] the econometrician David Hendry,[14] Peter Spirtes,[15] Jon Williamson,[16] and Christopher Hitchcock.[17]

To see how invariance enters, turn back to my original defense that causal laws earn their citizenship in our ontology because they sort effective from ineffective strategies.

4. Ooops: a badly flawed argument

TIAA is right, I argued: P(long life/TIAA insurance purchasers) is not a reliable guide for predicting whether I would increase my chance of living longer should I buy TIAA insurance. Still, for decision making under uncertainty and for cost-benefit analysis, we need a probability, if not the conditional probability, then the probability of a counterfactual conditional. What should it be? I, and a number of others at the time (like Brian Skyrms),[18] recommended a method that social scientists use to find causes, a method at the heart of Patrick Suppes's probabilistic theory of causality and also of the currently fashionable Bayes-nets methods for causal inference. To test whether C causes E, look to see if C and E are correlated once we have stratified on the other causes that can produce E by different routes.[19]

Having a healthy life-style and basic good health causes longevity. Being a college professor does not cause longevity but it is correlated with these other two factors; *that's* why being a professor is correlated with longevity. Stratify on these other factors and the correlation disappears. Among those with a healthy life style and basic good health, being a professor is no longer correlated with living long; nor is it correlated among those with unhealthy life styles and poor basic health. That is the idea behind using, not the conditional probabilities, P(E/C) versus P(E/-C), as a test for the causal law "C causes E," but rather the *partial* conditional probabilities, P(E/C&K_i) versus P(E/-C&K_i), where the K_i are the various different strata in which factors that affect E by different routes are held fixed. This is sometimes called formula CC (Causal Connection):

(CC) P(E/C&K_i) > P(E/-C&K_i).[20]

What justifies this method for inferring general causal claims? Here is my answer, which I think is the best we will get. In many domains of enquiry—maybe most—if two factors are correlated, there should be an explanation and the explanation should be provided by causal laws. One explanation for a correlation between two factors is that one of these factors causes the other. If we can eliminate all other possible explanations, this must be the right one. Stratification promises to do just that, by eliminating background correlations with confounding causal factors, such as healthy lifestyle and good diet.[21]

This formula might be a good idea then, though it does not provide a reductive account of causality, which many have hoped for. This was the core of my argument that we need some robust non-Humean concept of causality. We may start out trying to admit only associations. But

> Premise 1. For predicting the effectiveness of strategies, some associations are better than others.

Which are the good ones?

> Premise 2. If A is to be an effective strategy for R it is not correlation between A and R that matters but rather P(R/A&K_i) > P(R/-A&K_i).

But what is K_i?

> Premise 3. K_i describes a situation that is homogeneous with respect to all and only "other" causal factors for R. Nothing more nor less will do.

We may want to rely on associations but

> Conclusion. The difference between effective and ineffective strategies does not depend on associations alone but on facts about causal laws—on what the K_i's are.

The conclusion is correct. But my argument for it was no good. Look at premise 2. For an effective strategy for getting R's, I urged, we want what we would call a cause of R. Formula CC shows whether A is a cause of R. So CC shows that A is an effective strategy for R.[22] But even granting that strategies must be causes, which I will soon challenge, the formula that shows that A is a cause of R need not be the right formula for telling whether A will produce R when we introduce it.

Here is one reason we might have thought it was. Some arrangement or other must obtain, willy-nilly, when we set about changing A. Call it K_i

and consider the following formula which is a trivial consequence of the probability calculus:

$$P(R/K_i) = P(R/A\&K_i)\ P(A/K_i) + P(R/\text{-}A\&K_i)\ (1 - P(A/K_i)).$$

Suppose $P(R/A\&K_i) > P(R/\text{-}A\&K_i)$. Then it seems that increasing $P(A/K_i)$ will increase $P(R/K_i)$. So in situation K_i, introducing more A's will yield more R's.

No; and for two reasons. First, when we introduce A we generally change other factors—not just ones causally "downstream" from A, but factors that have independent routes for affecting R. In that case we will no longer be in situation K_i but in some new situation K_j. Then what's relevant is whether $P(R/A\&K_j) > P(R/\text{-}A\&K_i)$. But that's not right, for the second reason. $P(A/K_i)$ cannot change without the basic probability measure changing. In the probability measure designated by *P*, $P(A/K_i)$ is as it is. To reflect a different set-up with a different frequency of A's, and perhaps of R's, we need a new probability measure, say *P´*. The effectiveness of A as a strategy for getting R's then depends on a better **S**trategy **E**ffectiveness formula:

$$\text{(SE)}\ P'(R/A\&L) > P(R/\text{-}A\&K_i)$$

where *P´* is the new probability measure and *L* is the new (causally homogeneous) situation that will obtain once we have fiddled with the frequency of A's. Note that L may not even involve the same set of causal factors that K does since when we fiddle with the probability, we can easily alter the causal laws as well, changing facts about what other factors affect R and how. Indeed, fiddling with the probability without fiddling with the causal laws may be impossible since, I hazard, it is the exact form of the causal laws that sets the probabilities.

Let us return now to (CC) and get clear about just when it can tell us about the results of changing A:

$$\text{(CC)}\ P(R/A\&K_i) > P(R/\text{-}A\&K_i).$$

CC guarantees that P(R) increases if we do A only if among the causes of R we change A (and its downstream effects) and only A, so that K_i stays the same. Lots of our interventions are not like that. The California class-size reduction program is a well-known example.[23] The desired outcome, increased reading scores, didn't result. That's because the program was rolled out statewide over a short time. Halving the size of classes means doubling the number of teachers and the number of classrooms. There weren't that many qualified teachers available in that short a time, nor classroom space. The good effects of smaller classes were offset by the bad effects of poor teachers and poor facilities. (And, naturally, the most disadvantaged schools by and large got the worst teachers.)

Though this problem is serious, it is not the one I want to focus on. My worries here center on three further restrictions:

i) It's not enough that the factors in K_i stay fixed. It is because K_i picks out the other causal factors for R that the correlation between A and R with K_i fixed shows that A causes R. This means that in changing P(A) we must not change the other causal laws that have R in their consequent, neither adding nor subtracting from them.

ii) Nor must we undermine the very causal law that connects A and R if we are to hope to increase P(R) by increasing P(A).

iii) $P(R/A\&K_i)$ and $P(R/\text{-}A\&K_i)$ must also stay fixed, which may or may not follow from i) and ii) depending on views about whether the causal laws fix the associated probabilities.

These restrictions are just what we need to tell whether A and R are connected by causal law. But they are not tailored to tell what will happen if we introduce A to bring about R. So this is the first point I want to make in arguing that causal laws are not all that good as guides to effective strategies. My own way of making them so those many years ago rested on a mistake. CC is a good way of establishing or testing causal laws; it may even be a characterizing feature of them: Jiggle a cause *in just the right way* and the effect jiggles too. But *the right way* of jiggling A to test if it causes R is almost never the way we would expect to do it to produce R.

5. The bad penny

This mistake, like a bad penny, keeps turning up. Interventionist, invariance, and manipulationist accounts of causality abound now. These make the same connection between strategy and causal law that I did; and thirty years later many make exactly the same mistake. If A causes R is a true causal law, then CC holds for it. Hence we can increase P(R) by increasing P(A). This is at the heart of the manipulation theorems of Peter Spirtes and Clark Glymour and their associates and is also at the core of Pearl's semantics for counterfactual conditionals, as well as Woodward's more informal approach. But what about i) - iii)? Here is where *mechanisms* enter. The Mechanics take causal laws to be like levers: Jiggling at the cause end produces changes at the effect end. This won't work if the causal laws are not invariant when we use them, as required by i)-iii). So the Mechanics temper the levers so they will not be brittle.

There are various views about tempering on offer:

1. Causal laws generally **are** tempered (even under real manipulations).
2. Don't **call** something a causal law unless it is tempered, even if it passes other standard tests for being a causal law.
3. There is always some way to use even a glass lever.

I don't hold much truck with 2. and 3. 2. only connects causal laws and effective strategies by fiat, and 3. doesn't justify the huge efforts we put into causal laws in the hope of putting them to practical use. What about 1.? Are causal laws generally invariant, or invariant enough under realistic manipulations? There are two major kinds of invariance I worry about.

Locality. Few—perhaps none—of the causal laws we rely on for our day-to-day strategies are universal, God-given laws. Causal laws are generally only locally true. They almost always depend on an underlying structure that gives rise to them and to which they are bound. Manipulating the same cause does not produce the same effect for other generating structures.

Consider the lever that constitutes my car's throttle. Pushing on that lever makes the object the lever sits on accelerate. This causal law is true in my car and probably in yours and it's typically invariant across cars and trucks and even tractors. And it is a law that supports a strategy that most of us depend on many times a day, day in and day out. But this law is not true for most levers. Take a random sample of levers in the world and almost none of them will make the object they sit on accelerate. This important law is local, local to specific kinds of structures that give rise to it. I call structures like this, "nomological machines." A nomological machine is a stable (enough) arrangement of components whose features acting in consort give rise to (relatively) stable input/output relations. Most, if not all, causal laws hold only locally, relative to a nomological machine that ensures them.

Fragility. The second kind of invariance I worry about is *fragility*. So many of the causal laws I know break, and so readily, when we try to use them. Winding up the toy soldier causes him to march—if only you don't wind too tightly. Typing on my computer keyboard causes writing on my screen—if only I don't knock a cup of coffee over onto the keyboard.

Are these cases typical? I should like to urge that they are not untypical; that we know that; and that we worry about it. Consider the famous Lucas critique in economics.[24] Nobel-prize-winning Chicago economist Robert Lucas insisted that we cannot rely on even fairly well-established economic causal regularities to set policy. That's because these regularities arise from a more basic structure of individual decisions depending on "tastes and technology" and constrained by custom, law, regulation, institutional structure and human psychology—that is, they depend on a nomological machine. When the government intervenes, Lucas argued, it often disturbs this underlying system and in a way that destroys the very causal law that the government might use to predict the effects of its interventions.

6. Invariance and its values

In an irenic spirit we might note that, for reliable strategy, we do not need to settle how prevalent invariance is. Invariance is enough to ensure effectiveness. So base your strategies only on causal laws that are invariant.

This irenic remark may be good for peace in the profession, but it does not support the special status of causal laws. Recall, casual laws are supposed to have a special connection with strategy that is lacking for mere associations. But "What's good for the goose is good for the gander" applies to causal laws and associations. Tempering is good for both. There is nothing special about causal laws here. As Sandra Mitchell urges,[25] *any* claim that continues to be true when you act will be good for predicting the results of that action. What matters is not the causality but the invariance. So if you *add* invariance as a requirement, you don't need a causal law to start with.

But, one may question, perhaps causal laws are the only things that are invariant? Last section I argued that it is far from the truth that *all* causal laws are invariant under typical manipulations we might make based on them. Still, perhaps it is *only* causal laws that are invariant. I think not at all so. The argument that only causal laws are invariant rests on the same mistake as the argument that all (or almost all) causal laws are invariant—the bad penny argument.

Recall: The appeal of the view that causal laws are invariant under strategy manipulation typically comes from mistakenly assimilating manipulations we make in order *to achieve effects* with those we make *to test causal laws*. These latter manipulations—those that test causal laws—must leave the causal laws invariant by definition. The test would be no test if it added or subtracted the very thing it was testing for. Similarly, the appeal of the view that only causal laws are invariant also comes from a blinkered view of manipulation.

The standard argument proceeds via a hasty generalization to the contrapositive. It is common to remark that some paradigmatic non-causal regularities are not invariant. This is generalized to: No non-causal laws are invariant. From which follows: If it is invariant, it is casual. The standard examples are of joint effects, say E_1 and E_2, of a common cause. These are regularly associated—so called "spurious correlation." Jiggling one, it is noted, will not jiggle the other.

Or will it?

What *is* the case is that if E_1 is jiggled in just the way it must be jiggled to test if E_1 causes E_2—with all three conditions i)-iii) satisfied—then E_2 will not change in train. But this is not the kind of manipulation we generally make. Among those we do commonly make, a large number leave spurious associations like this invariant. And many of these we rely on—and reasonably so—for strategy.

Consider. I was at Stanford when I married Stuart Hampshire, who was just retiring from Oxford. I worried: Stuart, do you really want to leave Oxford and take the job that Stanford is offering you? Stuart reasoned just the way TIAA advised me not to. "If I buy TIAA insurance, I'll live longer" he assured me. His reasoning was correct. The association between buying TIAA insurance and the chance of living longer would be invariant for him because of the way he would do it. He would do so by becoming a professor in California, freed from the miasma of the Thames Valley.

This is not a singular example. I maintain that for the most part our day-to-day strategies are underwritten by the input/output regularities of a nomological machine. The input/output regularities look like causal laws from the outside. But sometimes when we study the inner workings, we see that the variable we manipulate is only spuriously related to the output achieved. To steer an aircraft, or an electron beam, you keep a needle centered on a dial. The inner workings of the steering mechanism assure that (just as with Stuart and the insurance for U.S. college professors) you can't do that without—inside the machine—adjusting a common cause of both the needle position and the direction of the aircraft, or the electron beam.

Or consider...you want to eliminate some serious hard-to-observe internal effects of a disease. To do so, you try out a variety of treatments that make more superficial external symptoms disappear. As with Stuart and the insurance or centering the needle to steer, you do so in the hope that, with one of these treatments, the way you make the external symptoms disappear is by eliminating the disease that is the common cause of both the superficial symptoms and the serious problems.

In offering this kind of example, I do not want to quarrel with the common assumption that a feature of the world cannot change unless a cause of it changes. Rather, I quarrel with the dual assumptions that the manipulations we make in day-to-day strategizing are likê the surgical interventions that test in the ideal for causes; and that the regularities we rely on, and reasonably so, must be causal laws. I repeat: Any conditional that is invariant when we manipulate the antecedent in the way that we will in fact manipulate it is as good any other for predicting the effects of our manipulations. It needn't be—and often isn't—a casual law.

Indeed, we don't even need invariance. What we need is truth.[26] As Julian Reiss urges:[27] Just go for a right answer. We do not need to start with claims that have been true somewhere, then look for them to be true in the situation at hand. We need a prediction that **is** true under the manipulations we propose. This is very much the point of another Nobel-prize winning Chicago economist, Milton Friedman, in his defense of unrealistic models in economics.[28] It doesn't matter if the model does not resemble the target situation; what matters is that the conclusion of the model is true for the target.

We are back now at the starting problem. With respect to effective strategies, what is special about causal laws? If we maintain that not all causal laws support strategies, only those that are invariant, we've made no special place for causal laws. Any invariant claim will do. Similarly, if we maintain that only those causal laws that are true given the manipulations proposed support strategy, again we have found no special place for causal laws. Any true claim will do. Once we have either invariance or truth we don't need causality.

Is there nothing then that is special about causal laws? Let's think about Friedman. The objection we raise in my philosophy of economics classes to Friedman's unrealistic models is that they are not useful. It is not correct that all we need from the modeling exercise is true predictions. We need some way of judging before the fact, some *marker*, for which these are. That is the point of the realistic model: If the premises in the model are true in the target, then the conclusion should be, too. The marker for true conclusions is true premises, which can in principle be checked independently.

In our case support from a causal law was supposed to be the relevant marker for an effective strategy. But this marker, I have been arguing, isn't reliable enough. The Mechanics add invariance. Though invariance won't save causal laws it can at least do the important job of underwriting strategy. It is not as good as truth, of course, since there will be myriads of true predictions about what happens if a course of action is undertaken that do not fall into already existing patterns. Invariance shares with truth, however, the drawback that it is hard to recognize. And despite the recent tsunami of commitment to invariance, there has not been much advance on how to know when we have it. Pearl *supposes* we have invariant causal laws and builds his semantics for effectiveness counterfactuals from these. Statisticians and econometricians have a variety of methods for testing for invariance, and this is a hot topic at the moment, but they are extremely limited in their usefulness, for two reasons. First, their range of application is limited. Worse, they are testing post hoc, not predicting. But for betting on strategies we need, rather, good ways for telling before hand, for predicting the range of the causal law.

We can do better, much better, if we stop skating on the surface. Unless we are dealing with the most basic laws of fundamental physics, invariance is a derivative feature. It comes from somewhere else. My version of the somewhere else is the nomological machine. However exactly we conceive it, that's where we need to focus to identify and to understand effective strategies.

To summarize: for effective strategies,

> Invariance is good, though it does not salvage causal laws. Once we have invariance we don't need causality.
>
> Truth is better. It provides us with a larger set of effective strategies.
>
> But neither truth nor invariance wears its label on its sleeve. We need something we can identify reasonably, readily, and reliably.

7. Nomological machines

Finally we have arrived at Thesis 3: The laws that support our day-to-day strategies are for the most part the temporally asymmetric conditionally stable input/output relations arising from the successful operation of a nomological machine.

A nomological machine, recall, is a stable (enough) arrangement of components whose features acting in consort give rise to (relatively) stable input/output relations. Some nomological mechanics are fairly "watertight": In normal conditions there's a constrained set of inputs they admit and these fairly reliably lead to canonical outputs. Others are more porous: They more often admit inputs that disrupt their canonical input/output relations. Nomological machines are often constructed by us. For example, toasters, cars, the UK financial system, hospital emergency admissions, and referral systems,... They also commonly occur naturally, at all levels: seeds, the Suprachiasmatic nuclei that Bechtel studies, the human body, ecosystems, ...

Consider my toaster. Drop in a piece of bread and push the lever. This regularly causes a piece of nicely browned toast to pop up a minute and a half later.

- This sequence has all the earmarks of a causal law.
- And it is invariant
  - so long as we push the lever on the toaster and not on the floor of my car
  - so long as my toaster's design is intact.
- And it provides an effective strategy (one I rely on every morning for my breakfast).

Importantly, this causal law is local to the toaster and guaranteed by it.

Or consider the nasturtium seed. *The Royal Horticultural Society Encyclopaedia of Gardening* notes, "In order to germinate, a seed needs water, air, warmth, and, for some species, light."[29] It is a relatively reliable causal law that the right mix of water, air, warmth, and light produces nasturtium seedlings. The law is local to the seed—you won't get seedlings by planting and watering basketballs—and it is guaranteed by the structure of the seed. As with the toaster, this causal law provides a strategy we employ again and again.

There are two central reasons we can rely on nomological machines for producing predictable results:

1. First, we often can learn to recognize these machines and what it takes to obtain an effective strategy from them. In particular we learn to recognize

- The machine itself. Often there are good markers by which we do so –
  - Labels
  - Other visible signs.
- Its canonical input/output relations, via
  - The manual that comes with the machine
  - Trial and error (e.g., "There is no way of distinguishing seeds requiring dark conditions from those needing light by physical examination; if the dark/light requirement is unknown, initially sow the seed in the dark; if it fails to germinate...place it in the light"[30])
  - Scientific study (e.g., RCTs to establish the output of treatment inputs on human bodies)
  - ...
- What endangers stable operation.
  - Again, often there are markers or visible signs by which we recognize damage.

2. Second is the nature of the inputs. For many nomological machines—in particular the ones that supply the strategies we use—

- The input causes are ones we can (or can learn how to) manipulate
- And can reasonably reliably do so without breaking the machine.

So my claim is:

> Most of the day-to-day effective strategies we rely on are underwritten by lawlike regularities generated by nomological machines and that many nomological machines, and notably the ones we rely on, have characteristics that make them both *recognizable* and *usable* by us. If you want to know if a proposed strategy will be effective you should look for the nomological machine that provides a reliable association between the strategy and the desired result. And if you want to know the range over which the strategy is effective, there is no substitute for figuring out what machine generates the connection. The lever on the door handle or on the toaster can look exactly like the lever on the floor of my car. To know whether pushing on it will accelerate you out of the garage rather than provide you with the morning's breakfast, you've got to know that you are in the car.

8. Why it's not just longer and longer causal laws

As philosophers you will naturally have a question in mind: So...we are not to pay too much heed to causal laws but instead focus on nomological machines. But aren't the principles that govern the operations of nomological machines causal laws? I say "No."

Along with other Engineers, I picture two tiers at work in nature: nomological machines—engines—and the "invariant" laws they give rise to. You can try to flatten out this landscape by bringing the description of the nomological machine into the antecedent of the causal law it generates. Then it becomes a flat landscape with a big bump in it, and—you miss the point. Putting the description of the machine in the antecedent makes the causal laws very long, and very complicated; there must be an inexhaustible number of these laws; and there is no apparent system to them. More importantly for my thesis, this open-ended, higgledy-piggledy collection of causal laws is not God-given. These laws have to be made, and we have the power to make them, and to make them different, by building different nomological machines.

What matters about nomological machines, whether constructed from iron and steel or from flesh and blood, is that they are made of parts with features that have powers and potentialities. Because of these powers and potentialities the parts can be assembled, and reassembled, in different ways to generate different causal laws. To understand these causal laws, and in particular to understand the ways in which they are fragile and the range across which they can be relied on, you need to understand what the parts can do, how they can work together, and what enables and what endangers their working together.

How should we characterize the basic facts of nature that determine what things can do and how they can act in consort? I talk about powers, about enabling conditions, and about rules of combination.[31] I'd be happy for someone to do better. What matters is that these basic facts are not yet-more causal laws. And what matters for practice is that we cannot find out about them with the same techniques we use to discover causal laws.

This is the problem with evidence-based policy right now. There are a great many guides on offer about how to judge effectiveness, from the U.S. Dept. of Education, to the UK's National Institute for Clinical Excellence, to the International Association for Research on Cancer; and they are all much of a muchness. RCTs are their gold standard. But RCTs are just the bad penny yet again, and this time where it affects our lives. It is via principle (CC) that RCTs can establish causal laws;[32] and principle (CC), we have seen, can predict the effectiveness of policies only under the most constrained kinds of manipulations. The evidence-based policy fix is the same as that of the Mechanics: invariance, or external validity; use a causal law only where it will hold. And how are we supposed to know where that will be? The U.S. Dept. of Education is typical in its advice: You are entitled to assume a program will be effective in your school if it has passed two good RCTs in "a setting similar to that of your schools/ classrooms."[33] I advise instead that we try to understand the underlying arrangement that will enable the program to produce the desired result. And we won't do that by conducting more RCTs.

9. In Sum

My point in this address is to get you worried about causal laws. Because you should be. They are not all they are cracked up to be. I am particularly concerned about this because of my recent work on evidence-based policy,[34] which focuses exclusively on establishing causal laws. But causal laws are neither necessary nor sufficient for predicting the effectiveness of strategy. My second point is to get you worried about invariance as a fix. Thirdly, I want to get you to focus on nomological machines, which are the engines of orderly behavior.

We are worried about change. But we are far too conservative in our views about how to bring it about. We can jiggle causes to get the effects we want. Except that doesn't work unless the lever connecting the two is tempered. So we must be very cautious: Just jiggle very carefully, so as to leave the causal laws invariant. Tempering becomes a process that we cannot act without.

But cherishing invariance denies us a more powerful, more fundamental tool for bringing about what we want. Do not just try to jiggle the cause paying strict attention to leaving the causal laws as they are. Rather, change the causal laws. As the Mechanics claim, tempered causal laws are mechanisms, but in the nineteenth-century sense. They are repetitive, able to transmit causal influence, but without serious productive power. We can jiggle one end of a tempered lever to move things at the other; and that is a big help. But for fundamental change, we can redesign the engines.

My predecessor at the London School of Economics, Karl Popper, insisted that we should engage in only limited piecemeal social planning. That's precisely because in his view we can neither understand nor build big engines of change. The same premise counsels no planning at all, as with Robert Lucas. For if we don't understand the engines of change, we do not know the stretch of our inductions and even piecemeal social planning is dicey. But we can build toasters, and high speed computers, and artificial knees that work. In all these cases we create new causal laws and we do so by creating new nomological machines. What about schools, child welfare, and international HIV-AIDS policies then? Lucas is, unfortunately, all too right about many of the interventions we attempt. The causal laws we rely on are altogether too fragile.[35] And they are destined to remain so, I fear, so long as we hire Mechanics to devise them. Let's at least give it a go: Just say "No" to Popper and Lucas. With better science—and better philosophy to underwrite it—perhaps we can build new social engines that make for better societies.

Endnotes

1. Thanks to Sophia Efstathiou for a great deal of help in preparing the lecture slides and the written version of this address. Research for it was sponsored by the (UK) AHRC grant, "Contingency and Dissent in Science." My thinking on these issues has benefited greatly from working with both Julian Reiss and Damien Fennell.
2. Cartwright 1979, 420.
3. This claim—causal laws are not much of a special guide to action—may seem hopelessly vague. But I think it is the right form for the conclusion to take. It is, after all, the equally vague claim that they are a guide that underwrites the huge expenditures we make to find causal laws. It is also, I think, the right abstract formulation to sit atop a number of more concrete claims expressed more precisely. Exactly what these concrete claims say depends on what characteristics one takes causal laws to have, what kind of actions are in view, and what rules are proposed for predicting the results of actions from causal laws.
4. Bechtel 2008.
5. Cf. Craver and Darden (2005) and references therein for their work and that of others on mechanisms.
6. There are other senses as well. Notably, a call to provide the mechanism between C and E is often a request for the causal process that connects the two, but represented in a series of steps where each step is described in a way that brings it under a "familiar" causal principle. I think this is what John Pemberton labels a "causal roll forward model" (in "How-mechanisms," 2009, ms: John.m.pemberton@btinternet.com) and that Jeremy Howick talks about in discussing mechanisms in evidence-based medicine in Howick (forthcoming).
7. Wise 2007.
8. Wise 2007, 163.
9. Wise 2007, 166.

10. Wise 2007, 168-69.
11. Woodward 2003.
12. Pearl 2000.
13. Steel 2008.
14. Hendry 2001.
15. Spirtes, Glymour and Scheines 1993.
16. Williamson 2005.
17. Hitchcock 1993.
18. Skyrms 1980.
19. This isn't the formulation in any of the works at the time but it is the one I would choose now. Also note that the causal Markov condition of Bayes' nets stratifies on other "direct" causes of C, not of E. In both cases the effect is to control for common causes that may induce correlations between C and E.
20. If the notion of a causal route is taken as given, then K_i is a state description over the set containing one factor from each route into E that does not go through C. Others define it using the notion of a causal parent, but I think this latter notion is always representation dependent and hence does not provide a proper way to pick out the members of K_i.
21. Yet, common causes are not always the only other possible explanation. For instance, I'm always worried about what is called "selection bias" in picking out the populations whose probabilities are being studied. I also argue that non-deterministic causes can produce effects in tandem so stratifying on a probabilistic cause need not eliminate the correlation between its common effects. So looking for these kinds of partial correlations is a sure-fire way to test for causal laws only in special kinds of systems where the other possible explanations are already ruled out. And that's a problem since the formula is often put forward as if it were universally a good test, not the least by me in my earliest work.
22. For a single case one may require merely a cause-kind that is a cause on that particular occasion. I am here supposing that for deliberation we want a cause-kind that is connected by a causal law with the effect-kind.
23. Bohrnstedt and Stecher 2002.
24. Lucas 1981.
25. Mitchell 2003.
26. Or even just "truth enough for the purposes at hand."
27. Reiss 2009.
28. Friedman 1953.
29. Brickell 1993, 536.
30. *Ibid.*
31. Cf. Cartwright 1989, 1999, 2007a, 2007b.
32. For a defence of this see Cartwright 2007c, Cartwright (under review), and Cartwright and Munro (forthcoming).
33. U.S. Department of Education 2003.

34. See Nancy Cartwright's personal webpage for her papers on evidence. Available at: http://personal.lse.ac.uk/cartwrig/Default.htm
35. For a discussion of this with respect to HIV-AIDS polices, see Seckinelgin 2007.

References

Bechtel, W. 2008. *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. Routledge.

Bohrnstedt, George W. and Brian M. Stecher (eds.) 2002. *What We Have Learned about Class Size Reduction in California*. California Department of Education. Available online at www.classize.org, last accessed: 09/09/09.

Brickell, Christopher (ed.) 1993. *The Royal Horticultural Society Encyclopaedia of Gardening*. Dorling Kindersley Publishing.

Cartwright, Nancy and Eileen Munro (forthcoming). "The Limitations of RCTs in Predicting Effectiveness." *Journal for Evaluation of Clinical Practice*.

Cartwright, Nancy (under review). "Predicting 'It Will Work for Us': (Way) Beyond Statistics." In *Causality in the Sciences*, edited by Phyllis McKay Illari, Federica Russo and Jon Williamson. Oxford University Press.

——. 2007a. *Hunting Causes and Using Them: Approaches in Philosophy and Economics*. Cambridge University Press.

——. 2007b. *Causal Powers: What Are They? Why Do We Need Them? What Can and Cannot be Done with Them?* Damien Fennell (ed.), Contingency and Dissent in Science Project Discussion Paper Series, Technical Report 04/07, CPNSS, LSE.

——. 2007c. "Are RCTs the Gold Standard?" *BioSocieties* 2(2): 11-20. Also as Damien Fennell (ed.), Contingency and Dissent in Science Project Discussion Paper Series, Technical Report 01/07, CPNSS, LSE. Also in Nancy Cartwright's *Causal Powers* (2007b).

——. 1999. *The Dappled World: A Study of the Boundaries of Science*. Cambridge University Press.

——. 1989. *Nature's Capacities and their Measurement*. Oxford University Press.

——. 1979. "Causal Laws and Effective Strategies." *Nous* 13: 4, 419-37. Also published in Cartwright (1983), *How the Laws of Physics Lie*, Oxford University Press.

Craver, C. and Lindley Darden. 2005. "Introduction: Mechanisms Then and Now." Special Issue: Mechanisms in Biology, *Studies in History and Philosophy of Biological and Biomedical Sciences*.

Friedman, Michael. 1953. "The Methodology of Positive Economics." In *Essays in Positive Economics*. Chicago University Press.

Hendry, David. 2001. *Causality in Macroeconomics*. Cambridge University Press.

Hitchcock, Christopher. 1993. "A Generalized Probabilistic Theory of Causal Relevance." *Synthese* 97: 335-64.

Hoover, Kevin. 2001. *Causality in Macroeconomics*. Cambridge University Press.

Howick, Jeremy (forthcoming). *What on Earth was Medicine Based on Before Evidence-Based Medicine? A Philosophical Inquiry*. Wiley-Blackwell.

Lucas, Robert. 1981. "Economic Policy Evaluation: A Critique." In *Studies in Business Cycle Theory*. Basil Blackwell.

Mitchell, Sandra. 2003. *Biological Complexity and Integrative Pluralism*. Cambridge University Press.

Pearl, Judea. 2000. *Causality: Models, Reasoning and Inference*. Cambridge University Press.

Spirtes, Peter, Clark Glymour and Robert Scheines. 1993. *Causation, Prediction, and Search*. Springer-Verlag Lecture Notes in Statistics 81.

Reiss, Julian. 2007. *Error in Economics: The Methodology of Evidence-Based Economics*. Routledge.

——. 2009. "Causation in the Social Sciences: Evidence, Inference, and Purpose." *Philosophy of the Social Sciences* 39: 20-40.

Seckinelgin, Hakan. 2007. *International Politics of HIV/AIDS: Global Disease Local Pain*. Routledge.

Skyrms, Brian. 1980. *Causal Necessity: A Pragmatic Investigation of the Necessity of Laws*. Yale University Press.

Steel, Daniel. 2008. *Across the Boundaries: Extrapolation in Biology and Social Science*. Oxford University Press.

U.S. Department of Education. 2003. *Identifying and Implementing Educational Practices Supported by Rigorous Evidence: A User Friendly Guide*. Institute of Education Sciences, National Center for Education Evaluation and Regional Assistance. Available online at: http://ies.ed.gov/ncee/wwc/references/iDocViewer/Doc.aspx?docId=14&tocId=3 last accessed: 09/09/09.

Williamson, Jon. 2005. *Bayesian Nets and Causality: Philosophical and Computational Foundations*. Oxford University Press.

Wise, Norton. 2007. "The Gender of Automata in Victorian Britain." In *Genesis Redux: Essays in the History and Philosophy of Artificial Life*, edited by Jessica Riskin. University of Chicago Press.

Woodward, James. 2003. *Making Things Happen*. Oxford University Press.