Ceteris paribus laws need machines to generate them

John Pemberton (LSE)
Nancy Cartwright (Durham & UCSD)

May 2014

1. Introduction

This volume builds on and explores a long tradition that maintains that many of our useful and prized laws in science hold only ceteris paribus. The claim recorded in the usual rendering of such laws is true only under some special conditions, to which the cp-clause refers. The long-standing challenge is how to give meaning to the cp-clause without rendering the law claim vacuous. You will find many answers to this question in this volume[1]. There are also many who despair, arguing that the question cannot be answered and that we should give up on cp-laws[2].

Our contribution to an answer builds upon a long-standing claim of Cartwright: Most of the regularities that get represented as 'laws' in our sciences arise from, and are to be found regularly associated with, the successful operation of a nomological machine.[3] This tells us immediately information that must be referred to in the cp-clause of one of these laws if the entire cp-claim is to be true: whatever else must be added, we cannot do without reference to a nomological machine. We agree, for example, 'Ceteris paribus aspirins cure headaches', but insist that they can only do so when swallowed by someone with the right physiological makeup and a headache. So conditions strong enough to guarantee the successful operation of a nomological machine are a necessary component of a cp-clause if the entire law claim is to be true.

Besides providing a necessary condition on the truth of the cp-law claim, recognising the nomological machine has great practical importance. This is what we aim to make clear here. For instance, referring to the nomological machine makes explicit where the regularities are to be found, which is of central importance to the use of cp-laws for prediction and manipulation. Equally important, bringing the nomological machine to the fore brings into focus the make-up of the machine – its parts, their powers and their

---

[1] See e.g. the contributions from Andreas Hüttermann, Bernhard Nickel and Gerhard Schurz.

[2] See e.g. Schiffer 1991; Earman, Roberts & Smith 2002; Woodward 2002.

[3] See e.g. Cartwright 1989; Cartwright 1999.

arrangements - and its context case-by-case. We call this the ‘local ontology’ of the machine. The need to refer to what we call the nomological machine is widely recognised. For example, Donald Campbell and Julian Stanley’s famous text on methodology[4] counts the description of the machine among the interaction variables that constitute a threat to the ‘external validity’ or ‘generalisability’ of an experiment. These variables, they explain, ‘represent a potential specificity of the effects of X (the cause under test) to some undesirably limited set of conditions.’[5] Those defending the use of causal Bayes nets for causal inference and discovery also address the need to condition causal claims to what turn out to be specific generating structures by adding additional variables to refer to those structures. These additional variables can for many purposes do the job of getting the causal claims right – they hold only when these new interaction variables take on the right values. On the other hand, they divert attention from the detailed make-up of the generating structure – the ontology of the nomological machine. And an account of the ontology is just what matters for many scientific practices. For instance, recognising the ontology of the machine provides a practical and principled basis for identifying possible disturbing factors, which are the central concern of so much of the cp literature. In the remainder of this paper we shall show how in this way, and in many others, getting a grip on the machine ontology is central to the methods of science and their success.

To this end, in Section 2 we extend our previous explication of how regularities (recorded as cp-laws) arise from the repeated operation of nomological machines and show how some disturbances can be allowed for by extending the machine.

Section 3 then provides examples that show how scientific practice recognises the nomological machine and its ontology, and how it recognises how the regularities which are recorded as cp-laws arise from the machine. We also look at what happens when we try to use cp-laws without sufficient regard to the nomological machine from which they come, citing errors which occur in the misapplication of randomised control trials (RCT’s).

In Section 4 we show how the methods illustrated in Section 3 associate variabilities in the cp-regularities with aspects of the nomological machine ontology and its context. We show case-by-case how this knowledge allows science to use cp-laws as the basis of successful strategies.

Section 5 outlines some implications for the practical use of cp-laws.

Before we proceed we should note a few things we are not doing. First we are not offering truth conditions for claims of the form, ‘Ceteris paribus L’. Rather we make a claim about one thing that ‘ceteris paribus’ must generally refer to if such a cp-law is to be true: the successful operation of a nomological machine. So, for the most part, our account of the cp-clause should

---

[4] Campbell & Stanley 1963.

[5] Ibid p. 187.

supplement rather than conflict with other accounts. Second, we do not offer sufficient conditions the clause must refer to in order for the cp-claim to be true but only the general form of a necessary condition. We hope you will find in this volume further considerations that can allow the construction of a nearly sufficient condition. But we do want to stress that whatever is on offer, if it is not strong enough to imply the successful operation of a nomological machine, it is very likely not strong enough to render most of our law claims true. You might put our central point thus: Reference to a nomological machine is essential if the cp-clause is to render a law claim true and the understanding of just what the nomological machine is that renders the laws true is central to a wide range of scientific practices that put our law claims to effective use.

Also, some remarks on ‘nomological machine’. Two decades after Cartwright’s first work on nomological machines, mechanisms are now a widespread topic. In current philosophical discussion ‘mechanism’ has a great variety of meanings, from James Woodward’s invariant generalizations to step-by-step causal processes to mechanisms like those of William Bechtel, Peter Machamer, Lindley Darden, Carl Craver and Stuart Glennan, where the focus is on components and their organisation[6]. Readers more familiar with this later literature on organised component mechanisms can mentally substitute these for our ‘nomological machines’ for the purposes of this paper[7].

2. Nomological machines and the cp-laws to which they give rise

A nomological machine is a sufficiently stable arrangement whose features (e.g. parts, properties, powers) acting repeatedly in consort can give rise to processes and input / output relations sufficiently stable to be the subject of cp-laws[8]. For instance the Solar System, consisting of the Sun and the planets, gives rise to the elliptical orbits of the planets described in Kepler’s laws. In agriculture, the appropriate application of fertiliser to corn plants can increase the crop yield. And when I insert a coin in to the refectory vending machine, a can of drink arrives in the output bin shortly afterwards.

---

[6] For a useful review of organised component mechanisms see McKay Illari and Williamson, 2012.

[7] Nevertheless there are differences that matter for different purposes between the various types of organised component mechanisms and our nomological machines (as well as many important similarities) - we explore this topic further in a forthcoming paper.

[8] To aid readability, we avoid labouring distinctions between token and type level nomological machines (machine arrangements and change processes) wherever possible. To suppose a change process is repeatable is to suppose it to be type-level (and thus associated with a type level machine arrangement). Often a token machine can repeatedly instantiate a change process (e.g. a neuron depolarising, a cistern producing a flush) by supporting the repeated instantiation of a start arrangement. Sometime a nomological machine produces a characteristic change process at most once (e.g. a seed germinating, a nebula producing a star). See Cartwright & Pemberton 2012, Pemberton 2011.

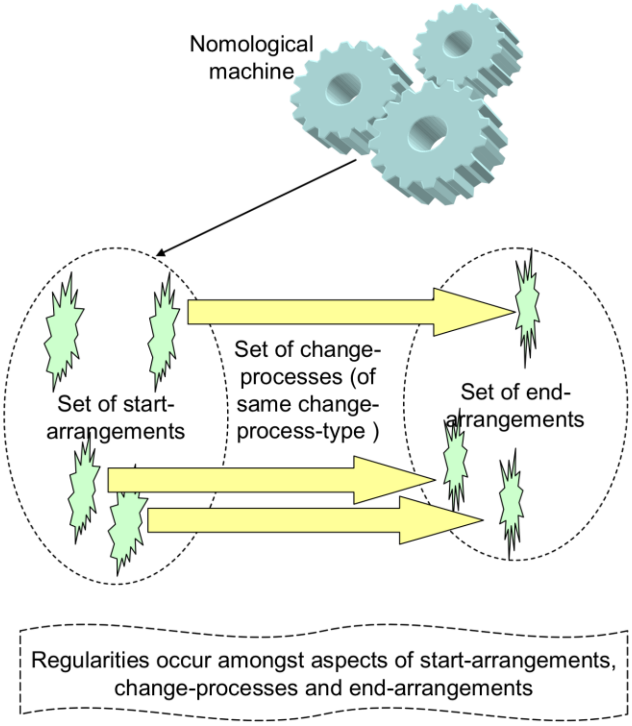

Figure 1

Some of the start-arrangements of the nomological machine give rise to change-processes of some characteristic type. The repeated operation of the machine can thus give rise to regularities amongst selected aspects of the start-arrangements, the change-processes and end-arrangements – regularities of the kind we record in cp-laws.

For example, Kepler's First Law tells us that ceteris paribus the path traced by a planet is an ellipse. Newton points to the nomological machine: the planet and its sun (i.e. the parts of the machine) gravitationally attracting each other on account of their masses, which gives rise to the repeated tracing of an ellipse by the planet (the characteristic change process).

Or consider the law that ceteris paribus aspirin cure headache. It is an uncontroversial claim of biological science that when this occurs, it is salicylic acid in the aspirin which migrates across the stomach wall in to the blood stream and then to the brain to act on relevant cells – an instance of this cp-law cannot occur without an aspirin entering the stomach of a suitable human body, i.e. the presence of the relevant nomological machine.

In practice the change which occurs on each run of the nomological machine is typically affected by factors external to the start arrangement of the machine: to 'disturbing influences'.

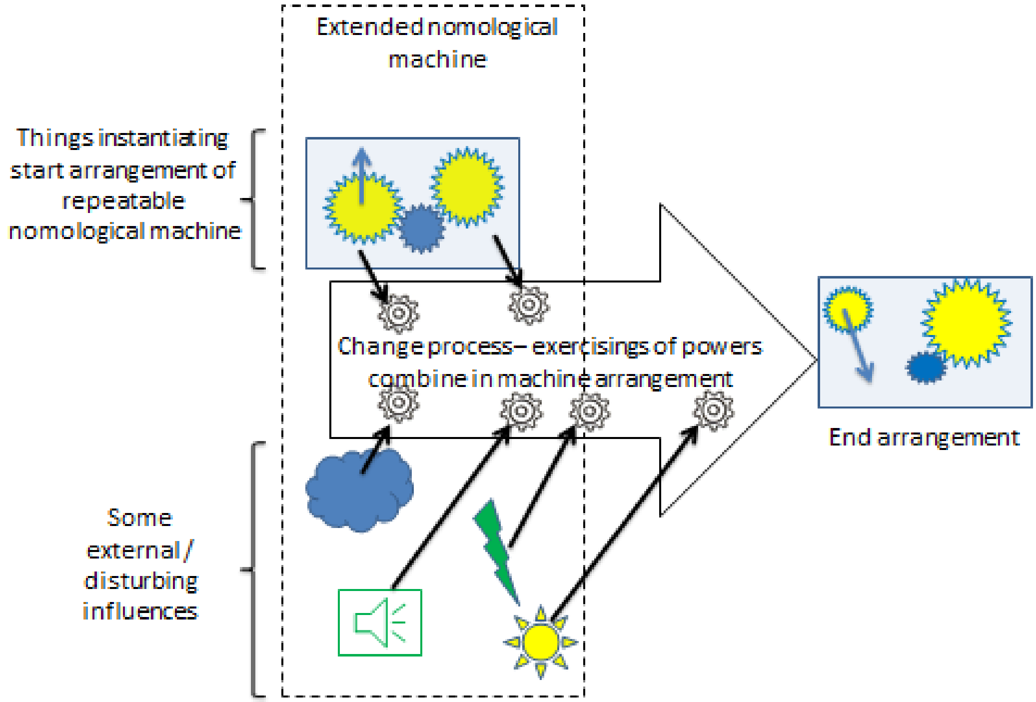

Figure 2

To deal with disturbances we can sometimes extend our characterisation to describe a new nomological machine. We include the disturbing influences as part of a revised start arrangement, and then consider the revised change process to which this revised arrangement gives rise. This extended machine may capture more accurately the change which occurs, but the increase in specificity of the arrangement can reduce the frequency of instantiation, perhaps considerably. In practice, then, such extended machines may have limited use. We typically focus on readily repeatable machines and must then manage the greater variabilities which occur due to disturbing influences. On occasions the disturbing influences may be so great that the characteristic change process of the nomological machine does not occur.

Alternatively we can make explicit cp-conditions which rule out factors that would disturb the machine significantly enough to disrupt the regularity reported in the law. For example in the case of Kepler's Law, we might extend our nomological machine to include other planets in the planetary system and perhaps even certain known comets. Doing so increases the accuracy of our predictive model, but also makes it considerably more complicated. Alternatively we might stick with Kepler's First Law and focus on the simple single planet nomological machine, explicitly ruling out other planets or comets in the cp-clause.

So in addition to what other papers in this volume note, the cp-clause does at least these two jobs: it conditions the regularity to an appropriate nomological machine and it rules out disturbances that would interfere with the machine in any way that would disrupt the regularity. These two conditions can be collapsed into one, so that we can say, for many if not all cp-laws, “ceteris paribus L” implies “L holds relative to the successful operation of a nomological machine appropriate to produce it.”

3. Using nomological machines and their cp-laws in science and technology

In this section we review some key methods used by science in support of its claims and practices. We show how these methods use nomological machines, and, in particular, how the uses made of the regularities (recorded as cp-laws) to which these machines give rise make good sense once their origin in the nomological machine is recognised.

In setting up a typical laboratory experiment (e.g. in physics, chemistry, biology), parts are brought together and arranged in the right way – placed in the right positions and connected in a prescribed manner: a nomological machine is constructed. We build a pendulum attached to a timing device, or rig up a neuron so that we can stimulate it to fire, or place chemicals in a test tube under sterile conditions. The experimenter’s next step is to get the machine working in the way anticipated: the pendulum swinging, the neuron firing in response to stimulus, the chemicals producing a characteristic reaction with the right sort of stages. Now she has a change process of the characteristic type, an instantiation of a repeatable experiment: the bob is stationary at the start with the string taut, it accelerates moving along an arc downwards and reaches its greatest speed at the bottom of its swing. The repeatability of the change process supports regularities between changes which occur during the change process. The release of the pendulum is followed by its movement through the low point of its swing; the stimulus of the neuron is followed by the firing of the neuron; the sparking of the chemicals is followed by some characteristic reaction.

Producing such a characteristic change process is often a starting point for further investigation. The experimenter may now vary aspects of the start arrangement or introduce some new influence during the running of the machine to see what happens. Typically the change which is made by the experimenter is carefully calibrated and is itself repeatable – we can think of the magnitude of this change as establishing a *setting*[9] for further runs of the machine[10]. Typically the experimenter explores the relationship between the setting and the effect by repeated running of the set-up with differing settings: e.g. between the length of the pendulum and its time period, the amount of tetanus applied to the neuron and its frequency of firing (and perhaps strength

---

[9] The setting is not a change event which forms part of the change process of the nomological machine, but may be a measure of some aspect of such an event or a measure of the magnitude of some cause of this change. (see below).

[10] Where we have a stable machine (e.g. a pendulum) a setting (e.g. the length of the pendulum) once changed may continue to apply to further runs of the machine until changed again.

of firing), the ratio of chemical inputs and the amount of carbon monoxide produced. Where a reliable relationship is established, we suppose the machine supports a functional relationship between the identified setting and the effect variables. The production of cp-regularities from the machine is central to establishing repeatability – and such repeatability of experiments is central to the methods of science.

The designer of a piece of technology, like the laboratory experimenter, typically brings parts together in some arrangement, perhaps building prototypes to see whether a new design does what is hoped. The design may be amended to increase speed and reliability, to reduce cost, to make manufacture easier, or to meet other objectives. This is an expert process centrally concerned with the machine as a whole.

However, many of the most common man-made machines are designed to be used by non-experts, people who need have no idea of the details of how the machine operates but who can simply follow easy instructions: push the button and wait, insert a slice of bread and depress the lever, insert a £1 coin in the slot. The human action brings about the start arrangement for the change process of the machine, or perhaps the first stage of change, so that the later change events of the change process now follow. The user thus exploits cp-regularities which arise from the change process the machine is designed to support – the action is followed shortly afterwards by some change event of the machine: the washing of my clothes, the popping up of a slice of toast, or the arrival of a drink in the output bin.

As well as supporting the cp-regularity between the depressing of its lever and its popping up of toast, toasters typically also provide a control which can be used to set the brownness of the toast. Manmade machines are typically designed to allow some settings to be easily controlled to produce outcomes relevant to the user. As non-expert users we learn how to adjust the settings by reading the manual. The regular association between setting and outcome, which can be recorded in a cp-law, is underpinned by the design of the machine, e.g. by the reliability of the parts and their secure attachment together so as to permit only well-defined movements of the parts. The law holds cp - relative to the machine and its proper working, and that fact is well known: We do not expect browned toast if we depress the lever on the toilet cistern nor if we have smashed the toaster with a hammer.

The cp-laws generated in our controlled experiments and by our technological devices come with the nomological machine already given. Generally we are not so lucky. We have a host of relatively well-established regularity claims that we would like to be able to rely on. But we know them to hold only cp: under some specific conditions or other. When are our predictions based on these regularities likely to be reliable? If our argument is right, these conditions will generally include the proper operation of a nomological machine. This, we claim, is reflected in scientific attitudes to the cp-laws and the predictions we might make from them.

Suppose that some of our experiments concerned with bacteria go wrong – the petri-dishes are not correctly sealed, or perhaps they are knocked, and the result is a blue-green mess of mould rather than the nice bacterial culture intended. Now we have noticed this phenomenon, we might tentatively posit a cp-law: when we have blue-green mould contaminating our petri-dish, certain bacteria in the dish die. Initially we may have little or no account of how the cp-regularity comes about – so how should we proceed? The next step in a typical scientific approach is to try to reproduce the phenomenon: perhaps a sample of the mould is isolated and then introduced into a petri-dish containing a suitable bacteria culture. Notice how this approach pre-supposes that the cp-law arises from a nomological machine – a machine which the scientist now tries to reproduce and run repeatedly in order to confirm the cp-law. And once this is achieved, the nomological machine is typically manipulated to explore its workings: Use differing moulds to see what types of mould work. Use differing bacteria to see which are affected. Investigate the changes which occur to the mould and the bacteria as the nomological machine runs, e.g. does the mould produce chemicals which we can isolate and show to be harmful to the bacteria?

Other cp-laws are further removed from the laboratory: The ancient Romans knew the cp-law that malaria occurs more frequently near swamps – they thought it was due to their production of foul air (hence the name). Effective action to control malaria only became possible when the responsible nomological machine was identified: the biting of a person by an infected mosquito. The proximity of swamps is only relevant as a breeding ground for insects. Spraying swamps may reduce the incidence of malaria, but only if we spray with the right thing - something that disrupts the relevant nomological machine, e.g. something that kills mosquitoes. Alternatively we might use mosquito nets or administer oral drugs to disrupt the machine. Here we intervene in the blousy world directly, but the methods we use still rely on our recognition of the nomological machine which underlies the cp-law.

CP-laws are also common within the bio-sciences. Recent studies of methods in this area (e.g. by Bechtel, Richardson, Darden and Craver) show that here too the response to a phenomena captured by a cp-law is an investigation which typically comprises the search for relevant parts and processes of change which underpin the regularity[11] – for the relevant nomological machine, we say. For example, when we investigate the firing of neurons, we discover that ceteris paribus the receipt of neurotransmitter particles is followed by a potential difference across the wall at one end of the neuron and then the movement of this potential difference along the neuron. The scientists investigating this cp-law discovered sodium selective pores in the lining of the neuron wall which open and close to control the movement of a cloud of sodium ions into and out of the neuron, thus supporting the

---

[11] See for example Bechtel & Richardson's account of decomposition and localisation methods (Bechtel & Richardson, 2010) and the many accounts of the explanatory role of mechanisms (e.g. Machamer, Darden & Craver, 2000).

propagation of a potential difference along the neuron. Discovering the relevant parts and their arrangement is often intimately linked with identifying the details of the change processes which occur to that arrangement. In practice scientists typically bootstrap to develop their account of the nomological machine: they first identify some of the relevant parts and their type and some aspects of the change processes – and then they look for more parts of the right sort of type, or perhaps develop their account of the nature of the change process, proceeding step by step. The identified regularity (cp-law) is then associated with the operation of this nomological machine – so that here too the nomological machine forms part of the cp-clause.

The socio-economic world is also a rich source of cp-laws. Here we are often focused on types of policy intrusion which aim to improve socio-economic conditions. We seek cp-laws of the form: a policy intrusion of type P is followed ceteris paribus by a change in some target socio-economic condition, measured by variable E. On our account, such policy intrusions typically act to change some pre-existing nomological machine.

Consider the example of child malnutrition on the Indian subcontinent. We suppose that the nutritional condition of a child arises out of the family context which includes the members of the family, the wealth available, and the degree of commitment to the child's well-being and the understanding of nutrition within the family. This family context is the relevant nomological machine which determines the child's food intake during their early years and hence the child's nutritional condition. A policy intrusion trialled in Tamil Nadu involved giving financial assistance to mothers and educating them about child nutrition, e.g. teaching them about the relative benefits of available foodstuffs. As in typical well-run policy programmes, the policy intrusion was well-defined and carefully administered, and its effect, the level of malnutrition amongst the children of these families, carefully measured. On our account this intrusion represents a change in the start arrangement of some of the nomological machines (i.e. families with young children) which now run in modified form.

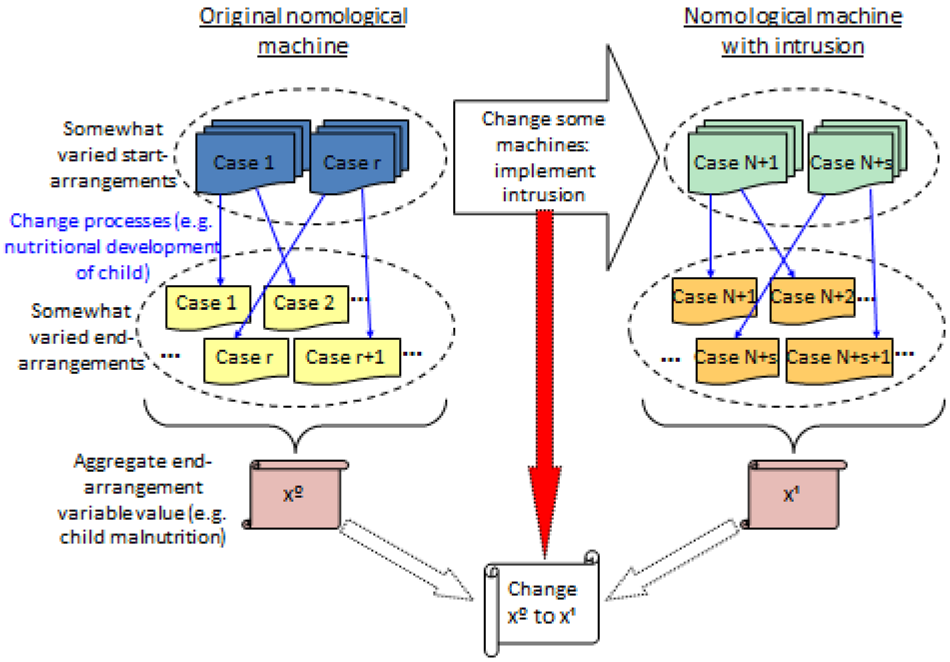

Figure 3

In this example, each of the cases in figure 3 represents a family with young children, either without policy intrusion (left-hand side) or with the intrusion, i.e. financial assistance and education (right-hand side). The results of the Tamil Nadu case showed a marked reduction in the level of child malnutrition – the policy achieved its aims.

On our account, possible intrusions depend on what amount to hypotheses about the nature of the nomological machine: its location, the nature of the parts and their arrangement, and the change process which characteristically occurs. In the best cases intrusions are designed using knowledge of how the nomological machine works. But what about when these are proposed without much knowledge of the underlying machine structure? In that case, we had better be very careful. Even if we are able to establish some regular association, even a regular causal connection, between something we do and some final state of interest, we must always remember that this regularity only holds ceteris paribus – relative to the successful operation of a properly operating nomological machine appropriate to generate it. So where we have little knowledge of the machine structure and operation, we have little knowledge of where that regularity will hold.

This is exactly the kind of problem that arises in the current craze for randomised control trials (RCTs) as evidence for genuine causal connections in evidence-based policy. The promise is that we can use RCTs to establish whether an intrusion genuinely causes a targeted effect, and that we can take this as good evidence that the intrusion will produce that effect if we adopt it in our setting. Whether it can serve as evidence will though, we know, depend on the machine structure in the RCT setting and the machine structure where the policy will be implemented. But that important fact is almost never acknowledged.

When the policy intrusions used successfully in Tamil Nadu were reproduced in Bangladesh they failed to improve child nutrition levels. Why? In Bangladesh it is the mother-in-law who is responsible for family finances and food buying and it is the men in the family who do the shopping, so that the education of the mother proved ineffective. In Bangladesh policymakers sought to make use of the cp-law used in Tamil Nadu without having regard to the nomological machine which gave rise to that law. When we consider the nomological machines in Bangladesh it is clear they differ in relevant ways from those in Tamil Nadu – we should not expect our original cp-law to hold there.

The RCT is now a popular form of test for causal relations in situations where we have little handle on what the underlying structure is like and what the other causes are that can affect the expected outcome given that structure. That's because RCTs can, at least in the ideal, allow us to establish casual connections between inputs and outputs with little knowledge of these other matters. Often an RCT is chosen to focus on a particular causal factor, for example the level of nitrates in the soil as a factor in crop yield. The RCT provides a tool to address the problem of possible natural co-variation of nitrates with other nutrients that might be responsible for improvements in yield: we suppose that if we control and vary randomly the amount of nitrates, we can test directly the impact of this factor separated out from the impact of any co-variates.

In general, randomisation and masking at all possible points in the experiment helps us to address the problems of co-variation, even where the relevant co-variates are unknown. So we are able to establish a causal principle - but of course a ceteris paribus causal principle. And, as we have been arguing, one of the central facts the cp-clause refers to is the structure and operation of the nomological machine that gives rise to it. So when the nomological machine is unknown, the RCT is a tool of very limited use. We can establish a ceteris paribus causal law but we have no basis for predicting where it will hold. This is particularly problematic for social policy since it is well known that the social structures that give rise to causal connections among socio-economic factors can vary dramatically even within small geographic regions. An action that is bound to cause offense given one social structure can be a required politeness in another.

In our view the RCT comes into its own not where we are hugely ignorant but rather where we already have a reasonable degree of knowledge about the structure of the machine and are probing for more information. For instance, as in Figure 3, varying the inputs or bits of the machine structure in controlled ways to see what the modified machine structure gives rise to, or to learn more about just what role that particular input or bit of structure plays in the overall operation.

In this type of research the analyst focuses on arrangements of things of specific types, looking for instance for those which can bring about the change process which is of interest: fecund families on the Indian subcontinent, plantings of corn seeds, humans satisfying certain health criteria. And particular attention is paid to the characteristic change processes which occur (e.g. child development, crop growth, disease development) and aspect of the outcome arrangements at suitable times (e.g. the nutritional health of children, crop yield, T-cell levels). How does the analyst know to choose these things and these aspects? Their framework of understanding concerns these things and how they work, i.e. the relevant nomological machine. Whether they articulate it this way or not, their research is developing further knowledge of the nomological machine responsible for the cp-regularities under study.

4. Managing variabilities in the machine regularities

Once a machine structure that gives rise to a cp-regularity has been identified, the sources of variability within the repeat operation of the machine can then be identified as well, thus linking variabilities in machine regularities to the ontology of the machine and its context. There may be variation in each of the following aspects of the start arrangement, which may influence the precise change which occurs:

- Machine parts. Although the things in the start arrangement must be of some prescribed type, with prescribed properties/powers, they will vary somewhat from case to case (e.g. differing £1 coins, with differing wear and tear, may be used in the machine).

- Arrangement of things. At the start the relevant things must be arranged in some prescribed layout, but the exact arrangement may vary and still produce the regularity reported in the cp law (e.g. differing precise positions of cans in the vending machine).

- Triggering change. Typically a change process is initiated by some triggering change. This triggering change may be the receipt of some input (e.g. a coin or neurotransmitter particles) or some rearrangement of the machine parts (e.g. the pushing of a button, the depressing of a lever). The precise quality (e.g. strength, speed, properties) of the change within the relevant trigger may vary (e.g. differing styles of inserting the £1).

Variabilities can also arise from factors external to the nomological machine, from ‘disturbing influences’. These are typically things in the vicinity of the machine which have the power to influence the relevant change process (e.g. things which jam the moving parts of the machine).

Many cp-laws are useful despite considerable variability, especially in the blousy everyday world, where there are typically very many disturbing influences. Where we have a usable regularity, these disturbances are either sufficiently infrequent or sufficiently small. Consider a machine designed to eject a cricket ball so as to give catching practice, with settings chosen so that the machine typically projects the ball over the boundary. On some occasion the ball may strike a hapless woodpigeon at mid-off, or be affected by the explosion of a nearby hand grenade, so that it does not reach the boundary – but such occasions are rare, we suppose. More typically, wind, rain, humidity, insects, etc. will vary the exact trajectory of the ball on each occasion but will not prevent the ball reaching the boundary. The cp-law that a ball flies off to the boundary when the button is pressed is useful despite such considerable variability.

However in science, we often seek to limit the degree of variability in our cp-regularities. Here we focus on circumstances where the disturbing influences are severely limited so as to permit the regularity of interest to obtain almost without exception: in the laboratory, in well-shielded machines, or perhaps in some astronomic or microscopic arrangements.

In the case of laboratory experiments, typical methods for managing variabilities include the isolation of the experimental set-up, using parts which are of known provenance so as to ensure their relevant quality and thus the reliability of their relevant behaviour, ensuring a suitably sterile environment, using measurement equipment which is well tested and calibrated elsewhere and has been found reliable, and using an expert laboratory technician with previous experience to ensure the arrangement is correctly set-up. Each of these methods is concerned with managing variations in the machine arrangement and its context of the types identified above.

Manmade machines designed for use in the world typically face considerable external variabilities – the design will typically include suitable shielding of parts, and the use of parts and methods of attachment which can withstand wear and tear. But consider our ball-throwing machine: we cannot shield its operation from the effect of local woodpigeons, wind, rain, humidity, or insects – the machine serves our purpose despite such variation. Nevertheless manmade machines are liable to break down sometimes – typically the characteristic change process is no longer produced or becomes too variable for the purpose at hand. Now we may ask an expert to mend the machine: typically we check the parts are still working individually and that they are still suitably attached – i.e. no loose screws or bolts. If the machine can’t be mended so that it again produces the required change process sufficiently reliably, then perhaps we can replace it with a new one.

Some natural nomological machines we can predict with great precision for long periods, such as the positions of the planets in the solar system, but more typically our ability to predict is limited: e.g. in meteorology, geology, economics, medicine. Although we may have some degree of confidence in the weather forecast a day ahead, and may find the forecast helpful, we may also make contingency arrangement for the occasions on which it is wrong: perhaps we take an umbrella. We know where earthquakes are more likely to occur and the typical range of their magnitudes, so here we design buildings to withstand shocks. Sometimes our predictive strategies are inadequate: a tidal wave damages a nuclear reactor – now perhaps we should review and change our strategies, e.g. build a flood barrier.

Statistical and econometric methods typically aim to identify systematic sources of variation – causal factors relating to some effect. But we do not identify all sources of variation: one way of managing the residual is to use probabilities. If the probabilities are sufficiently stable, and the residual variation is sufficiently limited, then our model may well be helpful for prediction, and perhaps for informing a strategy for control, always keeping in mind that what we are learning is tied to the nomological machine at work and does not apply at all where some different machine operates or where this one breaks down. So we always need to take care: our predictive mathematical models make strong implicit assumptions that the relevant nomological machine has stabilities which support the stable functional relations of the model and stable probability distributions. In general this is at best an approximation, and whether this approximation continues to hold during our forecast period is a matter we need to watch closely: if it starts to break down then perhaps we look for some change in the circumstances, perhaps some new technology or tastes, which could account for this change, and update our model.

In practice, knowledge of the type of circumstances in which a particular machine arrangement is instantiated, together with experience of the running of such machines on previous occasions, often allows identification and perhaps quantification of likely variabilities, and perhaps also the control of these variabilities as well as identification of where our predictions are totally irrelevant because some different machine is at work. Antibiotics may not reduce bacteria levels if they are consumed with high levels of alcohol – so doctors recommend not drinking when you are taking antibiotics. Educating mothers on nutrition is not effective in reducing child malnutrition if they don't have control of buying the family's food – such education programmes may not be recommended in cultures where the mother-in-law controls the household.

In general, our aim is not to eliminate variability (although the gravity probe experiment illustrates how far we can sometimes go[12]) but to recognize where the machine we have been studying will be found and to manage / control the variability sufficiently and to cope with the residual, so that our strategy is sufficiently effective in achieving its purpose. Where the relevant nomological machine is not operating or is not sufficiently stable in the right sort of ways, our methods are likely to be ineffective.

---

[12] See Cartwright, 1989, Section 2.4.2

5. In Sum

CP-laws arise from the operation of a nomological machine. So to use such laws to good effect, identify the relevant machine and its ontology, we enjoin. This is a method used by science with great success to predict and control. It provides a principled basis for recognising the relevant causal factors to use in our models and also for identifying the factors which must be held constant within our cp-conditions.

Because of the way they originate, CP-laws are both local and fragile: they hold just where and when the relevant machine is working correctly. The successful identification and use of such laws cannot be achieved by uncontextual general principles alone, but requires messy contextual knowledge, such as the knowledge of the laboratory technician or local anthropologist. Our review illustrates how science, in crafting its methods for explicating and using cp-laws, pays careful respect both to the nature of the nomological machine, and how regularities arise from its operation.

References

Bechtel, W. and Richardson, R.C. (2010). *Discovering complexity.* MIT Press.

Campbell, D. and Stanley, J. (1963). *Experimental and Quasi-Experimental Designs for Research on Teaching*. In: N.L. Gage Handbook of Research on Teaching. Chicago: Rand McNally: 171-246.

Cartwright, N. (1989). *Nature's capacities and their measurement*. Oxford: Clarendon Press.

Cartwright, N. (1999). *The dappled world: a study of the boundaries of science.* Cambridge: Cambridge University Press.

Cartwright, N. and Pemberton, J.M. (2012). *Aristotelian powers: without them, what would modern science do?* In Powers and capacities in philosophy: the new Aristotelianism. Edited by J. Greco and R. Groff. Routledge.

Earman, J., Roberts, J. and Smith, S. (2002). *Ceteris paribus lost.* In Ceteris paribus laws. Edited by J. Earman et al. Erkenntnis, 52 (Special issue): 281-301.

Hütterman, A. (this volume). *Ceteris paribus laws in physics*. Erkenntnis.

Machamer, P., Darden, L. and Craver, C.F. (2000), *Thinking about mechanisms.* Philosophy of science, 67 (March 2000).

McKay Illari, P. & Williamson. J (2012). *What is a mechanism? Thinking about mechanisms across the sciences.* European Journal of Philosophy of Science (2012) 2: 119-135.

Nickel, B. (this volume). *The role of kinds in the semantics of ceteris paribus laws.* Erkenntnis.

Pemberton, J.M. (2011). *Integrating mechanist and nomological machine ontologies to make sense of what-how-that evidence.* http://personal.lse.ac.uk/pemberto

Schiffer, S. (1991). *Ceteris paribus laws.* Mind, 100: 1-17.

Schurz, B. (this volume). *Ceteris paribus and ceteris rectis laws: content and causal role*. Erkenntnis.

Woodward, J. (2002). *There is no such thing as a ceteris paribus law.* In Ceteris paribus laws. Edited by J. Earman et al. Erkenntnis, 52 (Special issue): 303-328.