
The Limits of Exact Science, from Economics to Physics

Nancy Cartwright

*London School of Economics and
Political Science and University of
California, San Diego*

The idea of an exact science unified and complete has been advocated throughout the history of thought, but the sciences continue to cover only small patches of the world we live in. We may dream that the exact sciences will some day cover everything. But I argue that the very ways we do our exact sciences when they are most successfully done seems likely to confine them within limited domains. I discuss three cases to illustrate: the use of broad-scale non-experimental statistics for causal modelling across the social sciences, an economic model on skill-loss during unemployment, and the quantum theory of superconductivity. In all cases, where we can expect exact order depends on where we can fit our models. And by the nature of how models do-and should-get constructed in exact science, they fit readily onto only very special bits of the world around us. I also maintain that an ill-supported belief in the universality of our favourite exact science can lead us to adopt bad methodologies for carrying out the central aim of the sciences, namely to make the world the way it ought to be.

1. Introduction

The topic that I will focus on here is not the usual philosophical topic of realism in science—i.e., how accurately can the sciences represent the

Research for this paper was supported by the project “Modelling in Physics and Economics” at the Centre for Philosophy of Natural and Social Science at the London School of Economics. My thanks to all the members of that research group for their help with the ideas in this paper, but especially to Towfic Shomar, who designed the machine of section 2 and investigated a number of superconductivity examples with me, George Zouros, who has overseen the figures, Julian Reiss, who has commented in detail on the economics discussion, and Sang Wook Yi, with whom I have worked through the condensed matter physics and the philosophical arguments for section 4. Mary Morgan taught me the important lessons of the last paragraph of this paper.

Perspectives on Science 1999, vol. 7, no. 3
©2000 by The Massachusetts Institute of Technology

world, but rather the question of the range of the sciences—how much of the world can they represent. I am particularly concerned here with *exact science*: sciences that aim for formal theories, mathematical methods and precise prediction, with all the rewards that are expected to ensue, both in terms of precise control of the empirical world and of the relative prestige, power and money that is accorded the exact sciences and their practitioners.

The idea of an exact science unified and complete has been advocated throughout the history of thought, but the sciences have steadfastly resisted this kind of regimentation. They remain distinct and confined, and they continue to cover only very small patches of the world we live in. We may dream that the exact sciences will some day cover everything, but, I shall argue, that is not likely to be a dream that is even ‘in principle’ achievable. The very ways we do our sciences when they are most successfully done seems likely to confine them within limited domains. This is true, I believe, from the social sciences through the natural sciences and especially of fields like economics and physics that have embraced mathematics as their principal form of representation.

My thesis here is not meant to extend to any possible kinds of carefully controlled methods for gaining knowledge, but rather to special kinds of science, those that aspire to the positivist idea of exact science: science as a body of explicit knowledge, systematically organised, from which precise and unambiguous claims can be rigorously derived. It is my underlying view that it is this quite reasonable demand that scientific claims be precise and unambiguous that imposes limits on how far the sciences can stretch, for not much of the world lends itself to this kind of description. Basically the world as it comes, unengineered by us, is both messy and arbitrary and not the sort of thing about which the kind of knowledge we call scientific is possible.¹

I shall discuss three cases to illustrate:

1. The use of broad-scale non-experimental statistics for causal modelling across the social sciences.
2. An economic exercise from the work of one of my colleagues at the London School of Economics.
3. Quantum theory and in particular the quantum theory of superconductivity.

The role of models is key to my arguments. In all three cases, our scientific treatments require a model to represent the world, and the demands on these models constrain the extent of our treatments.

1. For an extended defense of this claim, see Cartwright (1999a).

2. Social statistics²

We use non-experimental statistics throughout the behavioural sciences as a tool for causal inference. Methodologies for doing so were explicitly developed by Sewall Wright in his work on path analysis; by the founders of econometrics and their immediate descendants in the early days of the Cowles Commission; and recently at UCLA and at Carnegie Mellon University, where groups working with Judea Pearl³ and with Clark Glymour⁴ have produced the most powerful and philosophically well-grounded techniques available, using directed acyclic graphs (DAGS).

Consider two examples, one from me and one from Judea Pearl. Pearl's example is shown in Figure 1, mine in Figure 2. Both examples deal with a number of measurable quantities that are represented as random variables and over which a probability measure P is defined. The graphs represent generic-level causal relations among the quantities. We assume that these two kinds of relations among the variables—probabilities on the one hand and causal connections on the other—are not independent, but that they constrain each other. Pearl and Glymour propose some general minimal constraints that they assume will always hold between the two. From that they have developed powerful techniques for inferring information about one or the other from partial information on the two of them.

Where do these two sets of relations come from? What gives rise to the causal laws and the probabilities represented here? For many empiricists this question will not make sense: causal laws and probabilistic—or, in the limiting case, deterministic—relations are fundamental to nature. What happens, happens on account of them; they do not hold on account of anything else. Well, I don't know anything about the growth of either eel worms or oats, nor about chemical fumigants and the possible interactions among the three that are depicted by Pearl. So I cannot postulate where his graph comes from. But I can tell you where mine comes from.

It comes from the machine in Figure 3. The machine gives rise to the causal relations and regularities depicted in the associated graph. The machine pictured in Figure 3 has fixed components with known capacities arranged in a stable configuration. It is the fixed arrangement among the parts that ensures both that the causal relations among the vertices of the consequent graph are as they are and that it makes sense to write down a probability measure over them.

I call machines like this 'nomological machines.' I have been spending a great deal of time lately looking at our basic theoretical claims in various

2. For a more detailed discussion of issues raised in this section see Cartwright (1997*a*) and Cartwright (1997*b*).

3. *Cf.* Pearl (1993).

4. Spirtes, Glymour and Scheines (1993).

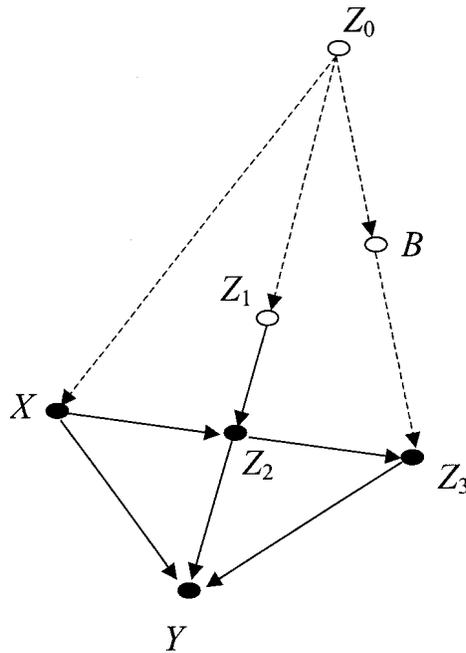


Figure 1: A causal diagram representing the effect of fumigants, X, on yields, Y.

- X: fumigants
- Y: yields
- B: the population of birds and other predators
- Z₀: last year's eelworm population
- Z₁: eelworm population before treatment
- Z₂: eelworm population after treatment
- Z₃: eelworm population at the end of the season

(Source: Pearl, J. 1995, "Causal Diagrams for Experimental Research", *Biometrika* 82: 669–710, p. 670)

exact sciences, especially in economics and in physics, at how we use them and at the kinds of evidence we can offer to support them. One of my conclusions from this work is that we always need a nomological machine to get laws—any laws, causal or otherwise. Sometimes God supplies the arrangements, as in the planetary systems, but very often we must supply them ourselves, in courtrooms and churches, institutions and factories. Where they do not exist, there is no sense in trying to pick out event-types

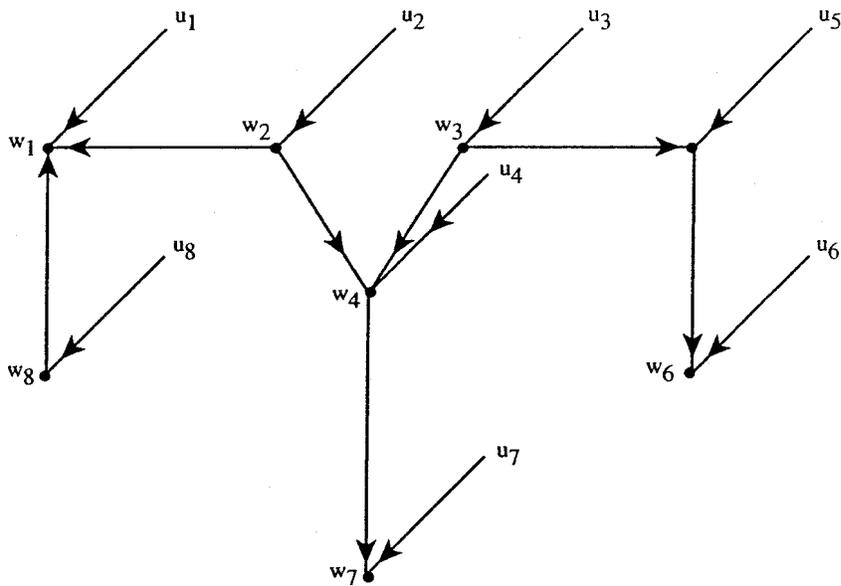


Figure 2: (Source: Designed by Towfic Shomar).

and to ask about their nomological relations, whether these nomological relations be causal laws or probabilities. That would be like asking: “Given we drop six balls in a bucket, what is the probability that a second bucket ten feet away will rise by six inches?” That question makes no sense unless we have in mind some background machine to which the buckets are bolted.

I can put this point in a more familiar language: probabilities are tied to chance set-ups. Chance set-ups give rise to the kinds and patterns of frequencies that we represent with probabilities; conversely, probabilities are the expression of the structures of the chance set-ups that generate them.⁵ A chance set-up may occur naturally or it may be artificially constructed, either deliberately or by accident. In any case probabilities are generated by chance set-ups, and their characterisation necessarily refers back to the chance set-up that gives rise to them. We can make sense of the probability of drawing two red balls in a row from an urn of a certain composition with replacement. But we cannot make sense of the probability of 6% inflation in the UK next year without an implicit reference to a specific social and institutional structure that will serve as the chance set-up that generates this probability.

The originators of social statistics followed this pattern, and I think

5. For a further discussion of this, see Hacking (1965).

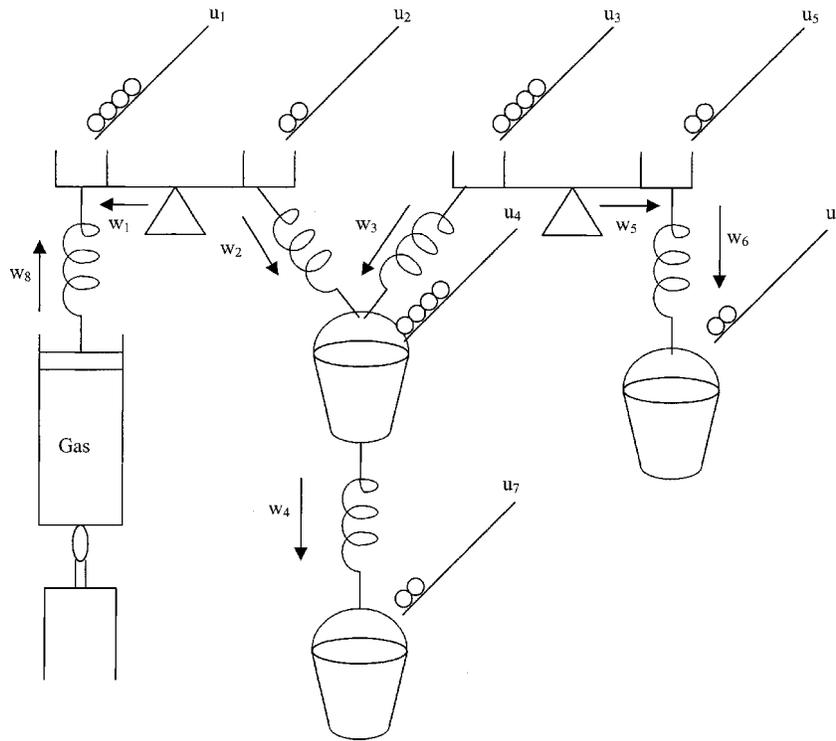


Figure 3: (Design: Towfic Showmar; recreation: George Zouros).

rightly so. When they talked about the ‘iron law of probability’ that dictated a fixed number of suicides in Paris every year or a certain rising rate of crime, this was not conceived as an association laid down by natural law as they (though not I) conceived the association between force and acceleration, but rather as an association generated by particular social and economic structures and susceptible to change by change in these structures. The same, I claim, is true of all our laws, whether we take them to be iron—the typical attitude towards the laws of physics—or of a more flexible material, as in biology, economics or psychology. I repeat the lesson about the dual nature of frequencies and propensities: probabilities make sense only relative to the chance set-up that generates them, and that is equally true whether the chance set-up is a radio-active nucleus or a socio-economic machine.

With reference to the behavioural sciences, we can summarise my point this way: sets of probability and causal laws like those depicted in Figures 1 and 2 cannot be free-standing. Rather they need a socio-economic ma-

chine to generate them. The view of course only matters if the kind of knowledge that we need to understand the operation of a socio-economic machine is not itself more knowledge of ‘deeper’ probabilistic and causal laws. That is just the claim I have long been arguing: the knowledge we need here is knowledge not of laws but of capacities (Cartwright 1989). And I claim that in fact much of the knowledge we have in exact science is just that—although, out of mistaken empiricist loyalties, we often try to reconstruct this knowledge as knowledge of laws.

In the simplest case these capacities can be thought of as probabilistic propensities. Patricia Cheng, a psychologist from UCLA, argues that in judging what causal relations hold given some information about associations, people “interpret their observations of covariations as manifestations of the operation of unobservable causal powers, with the tacit goal of estimating the magnitude of those powers” (Cheng 1997). What Patricia Cheng says people do in ordinary life is very close to what I see done in our most advanced sciences.

Rather than take you through a philosophical discussion of causal powers, I would just like to illustrate one case of where these philosophical matters make a real difference. The hero of the case is Amartya Sen. The case involves Sri Lanka, a very poor country which for years has had an active state entitlement system providing health care, education and food vouchers. Sri Lanka has also come out high on usual measures for standard of welfare among developing countries despite being in the low income group. This gives rise to an obvious question: is Sri Lanka’s social expenditure a significant cause of the high standard of welfare in its population?

Here is how some social scientists want to study this. You ask: is high social expenditure in general a cause of a high standard of welfare in developing countries? And to answer that you write down an equation like this:

$$H_{it} = a_i + bY_{it} + dE_{it} + l_i + u_{it}$$

where Y_{it} is per capita income; E_{it} is social welfare expenditure; a_i is a time-specific but country-invariant effect assumed to reflect technological advances; l_i is a country-specific and time-invariant ‘fixed effect’; d is the marginal impact of social expenditure on living standards; and u_{it} is a random error term. It is d that is the term of interest: the strategy is to try to use your social statistics to estimate it. If d is positive and large, the answer is ‘yes’; otherwise, ‘no.’

Now what strikes me is that this methodology is crazy. That is because what this equation represents is a ‘free-standing association’: there is no

good reason to think there is a chance set-up that generates it. You can of course use statistics as what the word originally meant—a *summary* of some data. In so far as there is data about the relevant quantities in a number of countries, you can write it out in lists or you can summarise it in statistics. And you can do that whether or not there is any probability over the quantities. But to go beyond that, to suppose that there really is some probability measure over welfare expenditure and welfare like that presupposed in the equation, you need a lot of good arguments, both at the phenomenological level—the frequencies pass strong tests for stability—and at the theoretical level—you have got some understanding of what kind of shared socio-economic structure is at work to generate them. The literature that I have reviewed does not provide the kind of substantive support needed on either front.

The equation I have written down comes from work by Anand and Kanbur (1995) criticising Amartya Sen for estimating the wrong equation, for estimating an equation in the level of the quantities rather than in their first differences. But this is odd since Sen does neither. Rather than the single graph pictured in Figure 4, which expresses a hypothesis Anand and Kanbur criticise about the causal relations that hold among the designated quantities across all developing countries, Sen's hypotheses dictate different graphs for different countries. And he does it this way exactly for the reasons I say. Each of the countries studied has a different socio-economic structure constituting a different socio-economic machine that will generate different causal relations true in that country and concomitantly different probability measures appropriate for the quantities appearing in these relations.

Look just at how Sen (1988) talks about South Korea and Taiwan on the one hand versus Sri Lanka on the other. In Taiwan and South Korea the immediate level of causation for poverty removal on Sen's account is employment expansion with growing demand for labour, and in this case export expansion was the immediate cause of the growing demand for labour. But, considering South Korea, "the fact that export expansion provided a very good vehicle for such explanation . . . has to be viewed in the perspective of the total economic picture." Sen tells us "*behind* [my emphasis] the immediate explanation" stand lots of other factors, like rapid capital accumulation, availability of a labour force suited by education to modern manufacture, and "overwhelming control of the organised banking sector" by the government. When these kinds of factors are arranged appropriately they constitute a kind of socio-economic machine in which causal links between export expansion and poverty removal can arise.

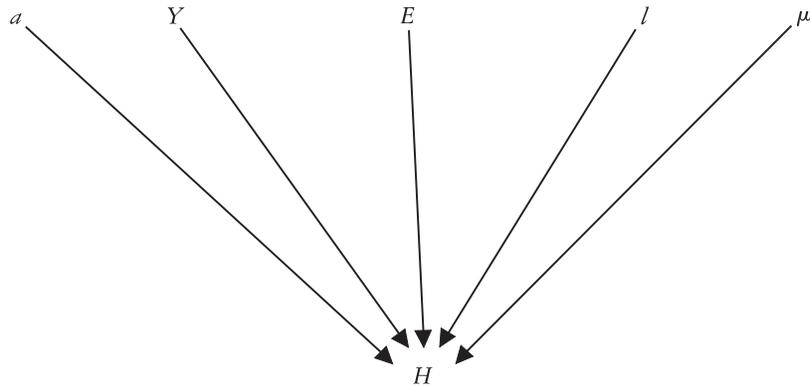


Figure 4: Anand-Kanbur Model.

$P(a, Y, E, l, \mu, H)$

a : technological advances
 Y : per capital income
 E : social welfare expenditure
 l : country-specific 'fixed effect'
 μ : 'error'
 H : living standard

3. An Economics Example⁶

The paper we will look at for an example from economics is titled "Loss of Skill during Unemployment and the Persistence of Unemployment Shocks" by Christopher Pissarides (1992). The idea investigated in the paper is that loss of skill during unemployment leads to less job creation by employers which leads to continuing unemployment. The method is to produce a model in which f_t , which represents the probability of a worker getting a job at period t ,

- (i) depends on the probability of getting a job at the previous period, f_{t-1} , (i.e. shows persistence) if there is skill loss during unemployment; and
- (ii) does not depend on f_{t-1} if not.

The model supposes that there is such a probability and puts a number of constraints on it in order to derive a further constraint on its dynamics:

6. For a more detailed discussion of the issues raised in this section, see Cartwright (1997c).

- (i) $\partial f_i / \partial f_{i-1} \neq 0$, given skill loss
- (ii) $\partial f_i / \partial f_{i-1} = 0$, given no loss of skill.

The point for us is to notice how finely tuned the details of the model plus the constraints on the probability must be in order to fix even a well-defined constraint on the dynamics of f_i , let alone to fix f_i itself.

The model is for two overlapping generations each in the job market for two periods only: at the end of each period every worker is, at least for the moment, unemployed. “Short-term unemployed” refers to ‘young’ workers just entering the job market at a given time with skills acquired through training plus those employed, and thus practising their skills, in the previous period; “long-term unemployed” refers to those from the older generation who were not employed in the previous period. The probability f_i of a worker getting a job in the between-period search depends critically on the matching technology χ , the number of times a job and worker meet and are matched so that a hire would take place if the job and the worker were both available. By assumption, χ at t is a determinate function of the number of jobs available at t (J_t) and the number of workers available at t ($2L$). Wages in the model are determined by a static Nash bargain. The central features of the first model are listed in Figure 5. Variations on the basic model that relax the assumptions that all workers search in the same way and thus have the same probability for a job match are developed in later sections of the paper.

The trick in the derivation is to get f_i to be a determinate function (via χ) of a product of f_i and f_{i-1} , where f_i itself is a function of J_t and L . In this model the product form in χ is achieved by getting J_t , which is determined by profits to be earned from offering a job, to depend on the product of f_i and f_{i-1} . This comes about because J_t depends on the probability of a job being filled by a short- (or long-) term unemployed worker, which in turn is equal to the probability of a worker being short- (or long-) term unemployed—into which f_{i-1} enters—times the probability of a short- or long-term unemployed worker getting a job, which is indifferently f_i for both.

The derivation of persistence or non-persistence in the rest of the models in the paper also depends on the fact that the relevant probability analogous to f_i in Model 1 is a function of the product of $f_i f_{i-1}$.

What lesson do I wish to draw from looking at Pissarides’ search model? Turn again to Figure 5. It takes a lot of assumptions to define this model and, as we have seen, the exact arrangement matters if consequences are to be fixed about whether there is persistence in the dynamics of unemployment probability or not. Those arrangements are clearly not enough to fix the exact nature of the persistence let alone the full probability itself. In Pissarides’ first model, where job openings are endoge-

Figure 5: Assumptions of Pissarides' 1992 model 1.

1. Discrete time.
2. Two overlapping generations.
 - a. Each of fixed size, L .
 - b. Each generation is in the job market exactly two periods.
3. Each job lasts one period only and must be refilled at the beginning of every period.
4. The number of jobs, J_t , available at beginning of period t is endogenous.
5. Workers in each of their two life periods are either employed or unemployed.
6.
 - a. Output for young workers and old, previously employed workers = 2.
 - b. Output for old, previously unemployed workers = $2y$, $0 < y < 1$. ($y < 1$ represents skill loss during unemployment.)
7. Unemployed workers have 0 output, no utility, no income. (This is relevant to calculating wages and profits.)
8. In each period all workers and some jobs are available for matching.
9. Each job must be matched at the beginning of a period to be filled in that period.
10. In each period workers and jobs meet at most one partner.
11. The number of matches between a job and a worker is designated by χ , where
 - a. χ is at least twice differentiable.
 - b. The first derivatives of χ are positive, the second, negative.
 - c. χ is homogeneous of degree 1.
 - d. $\chi(0, 2L) = \chi(J_t, 0) = 0$.
 - e. $\chi(J_t, 2L) = \max(J_t, 2L)$
12. There is a probability that a worker meets a job at the beginning of t , designated by f_t .
 - a. f_t does not depend on what a worker does nor on whether the worker is employed or unemployed.
 - b. f_t is a function only of J_t and L .
13. There is a probability that a job meets a worker at the beginning of period t .
 - a. This probability is independent of what jobs do.
 - b. This probability is a function only of J_t and L .
14. The cost of opening a job and securing the output as described in 6. is equal to $1/k$ (whether the job is filled or not).
15. Wages are determined by a Nash bargain.
16. Workers and employers maximise expected utility.

nous, the dependence of jobs on workers' histories must be engineered just so, so that J_i will be a function of the product $f_i f_{i-1}$. In the next model, where the product could not possibly enter through J_i , the facts about how workers search must be aligned just right to get the product into the search function S_i . And so forth.

My claim is that it takes hyperfine-tuning like this to get a precise probabilistic claim. As with the previous case, probabilities are associated with very special kinds of models before they are linked to the world. What I have pointed to—the need for a very finely honed model in order to derive any results—is not special to this case but is characteristic of economic theory. I think we can understand why.

The basic ideas I want to appeal to were formulated by Carl Menger (1883) in his attack on historical economics. Menger argued that economics could be a proper science and by 'proper science' he meant one that uses precise concepts that have exact deductive relations among them.

The paradigm is ' $F = ma$ '. In mechanics we get this exact relation, but the relation we have involves an abstract concept 'force,' whose relation to the world is always mediated via some more concrete concepts. These more concrete concepts are not geared to describe every arbitrary situation but are rather very specific in their form. The forms are given by familiar models, e.g., two compact masses separated by distance r , a linear harmonic oscillator, or a charge moving in a uniform magnetic field. Thus the concept force is limited in its application to situations that can be represented by these highly specialised models.

This is just what Menger said we should expect from economics: we can have concepts with exact deductive relations among them but those concepts will not be ones that represent features that are regularly found immediately in what he called "full empirical reality." This does not mean they never really occur, but if they do it would generally be in circumstances with special structures.

Much of the theorising we do in economics right now goes exactly the opposite direction to that recommended by Menger. It uses, not abstract or theoretical or newly invented concepts—like 'force' or 'energy' or 'electromagnetic field'—concepts partially defined by their deductive relations to other concepts, but rather, as Uskali Mäki teaches (Mäki 1996), very mundane concepts that are unmediated in their attachment to full empirical reality. We study, for instance, the capacity of skill loss during unemployment to produce persistence in employment shocks, or whether the limited enforceability of debt contracts will lead to inefficiencies, or whether when food shortages produce famines the effect is generally mediated by failures of food entitlement arrangements.

Nevertheless we want our treatments to be rigorous and our conclusions

to follow deductively. And the way you get deductivity when you do not have it in the concepts is to put enough of the right kind of structure into the model. That is the trick of building one of these economic models: you have to figure out some circumstances that are constrained in just the right way that results of interest can be derived deductively.

This makes for a kind of dilemma: on the one hand nothing follows from the concepts themselves without embedding them in a structure, and only special structures will yield any deductive consequences at all. But on the other, we don't want the circumstances really to matter. That is one of the chief reasons why economics theory is concerned with the robustness of the 'lessons' to be learned or understanding gained across models. But that is not what I want to focus on.

My point is that theories in physics and economics get into similar situations but via opposite strategies. In both cases we can derive consequences rigorously only in highly stylised models. But in the case of physics that is because we are working with abstract concepts that have considerable deductive power but whose application is limited by the range of the theory's models that provide more concrete forms for the abstract terms employed in the theory. In economics, by contrast, the concepts have a wide range of application but we can get deductive results only by locating them in special models.

4. The case of superconductivity⁷

We are invited to believe in the truth of our favourite explanatory theories because of their precision and their empirical successes. The Bardeen-Cooper-Schrieffer (BCS) account of superconductivity must be a paradigmatic case (Bardeen et al. 1957). We build real operating finely-tuned superconducting devices using the Ginzburg-Landau equations. And (with appropriate corrections!) we know that the Ginzburg-Landau equations can be derived from quantum mechanics or quantum field theory using the BCS model. So every time a SQUID (Superconducting Quantum Interference Device) detects a magnetic fluctuation we have reason to believe in quantum theory.

But the Hamiltonian used in the quantum equations of motion, like the classical force function, is abstract: we apply it to a situation only when that situation is deemed to satisfy certain other more concrete descriptions. These are the descriptions provided by what are called the 'interpretive' models of quantum mechanics. Albert Messiah's text *Quantum Mechanics* (Messiah 1961) provides four basic interpretive models: the central potential,

7. For a more detailed discussion of issues raised in this section, see Cartwright (1999b).

scattering, the Coulomb interaction and the harmonic oscillator, to which we should add the kinetic energy—which is taken for granted in his text.

Philosophers like C. G. Hempel and Ernst Nagel taught us to divide the principles of a theory into two kinds: internal principles, which give the relations between the abstract theoretical quantities, and bridge principles, which ‘interpret’ the theoretical concepts in more concrete terms. The quantum bridge principles give the corresponding Hamiltonians for each of the concrete interpretive models available in quantum mechanics. The point is: this is how Hamiltonians are assigned in a proper theoretical treatment; and in particular it is how they are assigned *in just those derivations that we take to be the best cases where predictive success argues for the truth of quantum theory*. When the Hamiltonians do not piggy-back on the specific concrete features ascribed to a situation—that is, when there is no bridge principle that licenses their application to the situation—then their introduction is *ad hoc* and the power of the derived prediction to confirm the theory is much reduced.

For the cases that give us really good reasons to believe in the truth of quantum theory, we need Hamiltonians that are assigned in a principled way; and for quantum mechanics as it is practised that means ones that are licensed by principles of the theory—by bridge principles. Now when we think about the range and limits of quantum theory there is something to be noticed. Bridge principles are just like internal principles in one respect: there are just a handful of them. And that is in keeping with the point of abstract theory as it is described by empiricists and rationalists alike. We aim to cover as wide a range as we can with as few principles as possible. But the fewer the bridge principles, the fewer the concrete models we have available to model the world. Since the theory can be applied only when its concrete interpretive models fit, the range of the theory will be severely restricted, even though the predictions within that range may be enormously precise and impressively accurate.

So far I have mentioned four basic bridge principles from Messiah. We may expect more to be added as we move from fundamental quantum theory to more specific theories for specific topics. The Bardeen-Cooper-Schrieffer account of superconductivity that I mentioned is a good example. I chose this example to study because it was one I knew something about from my work on SQUIDS at Stanford and from our research project on modelling at LSE. It turns out to be a startling confirmation of my point. The fundamental Hamiltonian of the BCS paper uses only the basic models I have already described plus just one that is new: the kinetic energy of moving particles, the harmonic oscillator, the Coulomb interaction, scattering between electrons with states of well-defined momentum,

and then in addition the ‘Bloch’ Hamiltonian for particles in a periodic potential (itself closely related to the central potential, which is already among the basic models). Superconductivity is a quantum phenomenon precisely because superconducting materials (at least low temperature, ‘type-I’ materials) can be represented by the special models that quantum theory supplies.

How much of the world altogether can be represented by these models is an open question. Not much, as the world presents itself, looks on the face of it like harmonic oscillators and Coulomb interactions between separated chunks of charge. Superconductivity is a case where a highly successful representation can be constructed from just the models quantum theory has to offer. My point is that with each new case it is an empirical question whether these models, or models from some other theory, or no models from any theory at all will fit. Quantum theory will apply to phenomena that these models can represent, and nothing in the theory, nor anything else we know about the structure of matter, tells us whether they can be forced to fit in a new case where they do not at first appearance do so.

5. Conclusion

So, how far does the kind of order that we represent in our exact sciences stretch? Consider physics with its highly abstract concepts. In so far as we are concerned with theories that are warranted by their empirical successes, the bridge principles of the theory will provide us with an explicit characterisation of its scope. The theory applies exactly as far as its interpretive models can stretch. Sticking to Messiah’s catalogue of bridge principles as an example, that means that quantum theory extends to all and only those situations that can be represented as composed of central potentials, scattering events, Coulomb interactions and harmonic oscillators.

The kind of economic theory of which we saw an example works in the reverse way. Economics uses very concrete concepts, like the level of employment or length of time between jobs, which do not have very strong deductive relations built into them, even given the handful of principles we may be willing to accept from economic theory. So a lot of very particular kinds of constraints need to be supplied by a model before any exact conclusions can be derived. And from my first case we learn that, where we want to postulate a determinate set of causal relations or a probability measure over a set of quantities of interest, there must be a socio-economic structure to generate them. We can understand many of our typical models in political economics as blueprints for these structures. In all cases, where we can expect exact order depends on where we can fit our models. And by the nature of how models do—and should—get

constructed in exact science, the kinds of models available are not very flexible. They fit readily onto only very special bits of the world around us. It seems to be in the nature of how we do exact science that what it can cover is very limited in scope.

The conclusions I draw are of course not certain. The kinds of models we build right now in our exact sciences may after all, despite appearances, really cover all the phenomena to which we might aspire to apply those sciences. Or we may develop different ways of doing exact science even more successfully that do not limit its application to situations covered by such restricted kinds of models.

Nevertheless, I maintain, the conclusions are well-supported by the evidence, and in particular they are at least as well supported as the hope that our exact sciences can take over all of nature. Still, both conclusions are clearly speculative; they go well beyond the evidence. As an empiricist I would thus urge that we make every effort to avoid having to bet on either. Why then am I at pains to argue that exact science has a good chance of being severely limited?

I do not argue from a concern that I or my colleagues may have mistaken beliefs about the scope of exact science, for I do not see what belief matters, *except for its effect on action*. In my view the central aim of the sciences is to make the world the way it ought to be. And I am worried that an ill-supported belief in the universality of our favourite exact science can lead us to adopt bad methodologies for carrying out this aim.

The damage from the drive to extend our favourite theories to treat new problems of new kinds is not done when there is good empirical evidence for the promise of the proposed approach to solve those specific problems. Attempts at take-overs are not out of place when we have good reason to think the methods or ideas we have developed to solve one kind of problem can be employed to solve a different kind, when we can argue the case and make a reasonable bet, where the costs have been thought through and our assessments of the chances for success can warrant this way of proceeding over others. If it is not even objectionable to invest in the attempt out of sheer passion for order and rationality, nor to try our favoured theory just because we have no better ideas or because that is what we are good at. It is not objectionable so long as we are clear about what we are doing and what it is going to cost us, and we are able to pay the price.

But we are often not so clear. The yearning for 'the system' is a powerful one; the faith that our world must be rational, well ordered—and by the same principles—through and through, can often play a role where only evidence should matter. Our decisions are affected. After the evidence is in, theories that purport to be fundamental—to be able in principle to explain everything of a certain kind—often gain additional credibility just for that

reason itself. They get an extra dollop of support beyond anything they have earned by their empirical success or the empirically warranted promise of their research programme for solving the problem at hand. We are all familiar with the take-over attempts of strings and symmetries and dualities, of fundamental particles in physics, of quantum mechanics over classical, of Darwinism in social psychology and of equilibrium theory, rational expectations and game theory in political economy. In medicine one of the primary take-over theories is genetics. Let's consider it.

Genetics has proved very useful across a variety of problems. Genetic information about the heritability of the production of Down's syndrome babies has led to much better targeting of *amniocentesis* among young mothers. Our knowledge that glycogen storage diseases are due to single point mutations allows us to administer the somewhat dangerous tests for these diseases only to babies with a family history. Now that we understand that the serious mental handicaps of phenylketonuria (PKN) are due to a single point mutation that leads to too much accumulation of one amino acid and not enough of another with a resulting failure of neurological development, we can adjust the diet of the children affected till the period of development is over. Some even hope that we will learn how to use viruses as vectors to resupply the genetic material.⁸

For many diseases, however, other approaches may be just as, or even more, fruitful. For example, in breast cancer there is very substantial evidence that endogenous oestrogen levels are the major determining factor in the occurrence of the disease in the vast majority of cases. It is well known that endogenous oestrogen levels are affected by lifestyle, but little emphasis is put on this aspect of prevention: on finding out how diet, physical activity (exercise and work) and other modifiable factors may lower endogenous oestrogen levels, or on finding 'synthetic' approaches that could achieve the same end. A rational approach to the problem of breast cancer prevention would put as much emphasis here as on the genetic aspects.⁹

Is the level of effort and funding that goes into the gene programme versus the others warranted by the promise of these programmes for understanding and controlling breast cancer or does the gene programme get a substantial edge because it is the gene programme; because it is our best shot right now at a theory of everything? I care about our ill-supported beliefs that nature is governed by some universal theories because I am afraid that women around the world are dying of breast cancer when they need not do so because other programmes with good empirical support for

8. Thanks to Jennifer Worrall for help with these examples.

9. Thanks to Malcolm Pike for help with this example.

their proposals are ignored or underfunded. On the side of political economy, Chicago economists carry the conclusions of the putatively all-encompassing theory of rational expectations far, far away from the models where they can derive them to admonish the government against acting to improve social welfare; and the World Bank and the International Monetary Fund use statistical models of the kind I criticise in section 2 to discourage Third World countries from direct welfare expenditure. That is what is wrong with faith in the limitless theory.

My own research right now is not primarily concerned with economics or with physics or with any other single discipline from within. It is concerned rather with how to get the most out of our scientific knowledge as a whole. How do we best put together different levels and different kinds of knowledge from different fields to solve real world problems, the bulk of which do not fall in any one domain of any one theory? Within each of the disciplines separately, both pure and applied, we find well developed, detailed methodologies both for judging claims to knowledge and for putting them to use. But we have no articulated methodologies for interdisciplinary work, not even anything so vague and general as the filtered-down versions of good scientific method that we are taught at school. To me this is the great challenge that now faces philosophy of science: to develop methodologies, not for life in the laboratory where conditions can be set as one likes, but methodologies for life in the messy world that we inevitably inhabit.

References

- Anand, S. and Kanbur, R. 1995. "Public Policy and Basic Needs Provision: Intervention and Achievement in Sri Lanka." *The Political Economy of Hunger: Selected Essays*. Edited by J. Dreze, A. Sen and A. Hussain. Oxford: Oxford University Press.
- Bardeen, J., Cooper, L. N. and Schrieffer, J. R. 1957. "Theory of Superconductivity." *Physical Review* 108:1175–1204.
- Cartwright, N. 1989. *Nature's Capacities and their Measurement*. Oxford: Oxford University Press.
- . 1997a. "What is a Causal Structure?" In *Causality in Crisis? Statistical Methods and the Search for Causal Knowledge in the Social Sciences*. Edited by V. R. McKim and S. P. Turner. Notre Dame: University of Notre Dame Press.
- . 1997b. "Where Do Laws of Nature Come From?" *Dialectica* 51: 65–78.
- . 1997c. "Models: Blueprints for Laws." *Philosophy of Science* 64: S292–S303.

- . 1999a. *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.
- . 1999b. “Models and the Limits of Theory: Quantum Hamiltonians and the BCS Model of Superconductivity.” In *Models as Mediators*. Edited by Morrison, M. and Morgan, M. Cambridge: Cambridge University Press.
- Cheng, P. W. 1997. “From Covariation to Causation: A Causal Power Theory.” *Psychological Review* 104: 367–405.
- Hacking, I. 1965. *The Logic of Statistical Inference*. Cambridge: Cambridge University Press.
- Pearl, J. 1993. “Aspects of graphical models connected with causality.” in *Proceedings of the 49th session of the International Statistical Institute*, Tome IV, Book I, Florence: Italy, pp. 391–401.
- Mäki, U. 1996. “Scientific Realism and Some Peculiarities of Economics.” In *Realism and Anti-Realism in the Philosophy of Science*. Edited by R. S. Cohen, R. Hilpinen and Q. Renzong. Dordrecht: Kluwer.
- Menger, C. [1883] 1963. *Untersuchungen bei die Methode der Sozialwissenschaften, und der politischen Oekonomie Insbesondere*, Leipzig: Duncker & Humblot, transl., *Problems of Economics and Sociology*, Urbana: University of Illinois Press.
- Messiah, A. 1961. *Quantum Mechanics*. Amsterdam: North-Holland.
- Pissarides, C. 1992. “Loss of Skill During Unemployment and the Persistence of Unemployment Shocks.” *Quarterly Journal of Economics* 107: 1371–1391.
- Sen, A. 1988. “Sri Lanka’s Achievements: How and When,” In *Rural Poverty in South East Asia*. Edited by T. N. Srinivasan and P. K. Bardham. New York: Columbia University Press.
- Spirtes, P., Glymour, C. and Scheines, R. 1993. *Causation, Prediction, and Search*, New York: Springer-Verlag.