Explanatory Mechanisms and the Aims of the Social Sciences

I

Social scientists pursue a wide variety of different ends with their model-building practices. To name but a few, data models are built for measuring complex social phenomena; forecasting models are built for predicting the future values of target variables of interest; explanatory models are built for gaining a deeper understanding of phenomena of interest; policy models are built for analysing the likely effects of interventions.

Among many philosophers and methodologists of the social sciences, however, there seems to be a tendency to focus on one (or a few) of the variety of different ends at the exclusion of others. For example, positivists in the philosophy of the social sciences tended to emphasise the importance of description and prediction, professing their belief that genuine understanding is impossible, while their opponents thought only understanding was important.

Generally speaking, there are four possible reasons for a methodologist to ignore an aim of the science he or she is concerned with. First, one may think that there is a hierarchy among the different aims and focus on that or those which one believes to be the more important, more ultimate end or ends. For instance, one may believe that the major purpose for prediction is the confirmation of explanatory theories and hence focus on explanation while considering prediction only in its relation to the more ultimate aim. Second, one may think that some aims are unattainable and consequently focus on those that one believes to be attainable. For example, one may regard successful prediction of social phenomena as impossible and therefore concentrate on the more promising explanatory endeavour. Third, one may regard an aim as both important as well as attainable, but as methodologically trivial. For example, one may regard the measurement of inflation as an important purpose of economics but deny that it involves any problems of particular methodological interest. Fourth, one may think that important classes of models serve a variety of different purposes at the same time. For example, if one believes that good models do or ought to represent laws of nature and one believes that knowledge of laws of nature allows us to explain past events and predict future events by the same token, one can safely focus on either of the two aims because the other will be served in tandem.

In this paper I discuss a relatively recent movement in the philosophy of the social sciences, which emphasises the importance of causal mechanisms. According to this movement, which I am going to call the "new mechanist perspective" or NMP, theoretical explanation of social phenomena is the only or ultimate aim of the social sciences and causal mechanisms play an essential role in theoretical explanation. Against these claims I argue that causal-mechanistic explanation is neither the only nor the most important or ultimate aim of the social sciences and that investigating causal mechanisms will not usually be the optimal strategy when other aims of the social sciences are concerned.

The next section introduces the movement, derives a more precise statement of NMP and gives a number of examples of causal-mechanistic explanations for social phenomena of interest. In the two sections that follow I go through the four possible reasons why a number of non-explanatory aims of the social sciences might safely be ignored and demonstrate that neither of the given arguments succeeds. In conclusion, I argue in favour of a tighter fit between methodology and purpose and, since there are a variety of different purposes with more or less equal rights, for a more pluralistic methodology of the social sciences.

II

Thinking about mechanisms has become increasingly popular in recent social science and its methodology. While the idea that it is necessary to provide a description of a mechanism of a certain kind in order to adequately deal with social events and phenomena has been an item on the neoclassical agenda in economics for quite some time, this basic picture has now been

adopted explicitly and in its own way in some heterodox schools of economics and in parts of sociology, social psychology and by a number of philosophers of social science. In particular I am thinking of the critical realism movement in economics, sociology and general social science (see, *e.g.*, Lawson 1997; Lawson 2003for economics, Brante 2001 for sociology and Sayer 2000 for social science), various accounts of social mechanisms in economics and sociology (see in particular the contributions to the volume Hedstroem and Swedberg 1999a) as well as contributions to the philosophy of social science (see, *e.g.*, Elster 1983; Elster 1985; Elster 1989; Little 1991; Little 1998). While the individual accounts differ in their details, they share a number of fundamental convictions. The first common idea is the emphasis on theoretical explanation as scientific virtue. Many proponents of the NMP appear to claim that theoretical explanation is the only aim of the social sciences. A somewhat weaker form of this claim is that theoretical explanation is the most important or ultimate aim of the social sciences. Consider the following statements:

> Sociology must seek to be an explanatory science, implying a deepening of its own theoretical knowledge. […] I propose that the ultimate goal of sociology is to identify social structures harbouring causal mechanisms that generate empirically observable effects. […] According to Bhaskar… the object of sociology is social relations, and the task is to explain the reproduction and transformation of social relations. (Brante 2001, pp. 168 and 178, emphasis original)

> In addition, the impossibility of engineering, and the absence of spontaneously occurring, closed social systems, necessitates a reliance upon non-predictive, purely explanatory, criteria of theory development and assessment in the social sciences. (Lawson 1997, p. 35)

> The main message of this book is that the advancement of social theory calls for an analytical approach that systematically seeks to explicate the social mechanisms that generate and explain observed associations between events. (Hedstroem and Swedberg 1999b, p. 1)

The second common feature is the idea that the social world is structured or layered in the sense that, on the one hand, there are higher-level (macro or aggregate or readily measurable) events and event-patterns, such as the rise of the dot.com industry, German inflation in 1923, the (in)effectiveness of crime prevention programmes, the effects of people's religious beliefs on their economic behaviour, and, on the other hand, there are the lower-level (micro or individual or theoretical) structures and processes that give rise to the higher-level events or phenomena. Let us call social event tokens and types of interest "phenomena" (basically adopting Ian Hacking's terminology, see Hacking 1983), the higher level the "empirical layer" and the lower level, the "underlying layer" (bearing in mind that the lower level can itself be structured so that there may be a hierarchy of ever deeper layers within the social world).

The third common feature is the thesis that the only or ultimate aim of social science is fulfilled if and only if the social phenomenon S (at the empirical layer) is explained in terms of the causal mechanism(s) (at the underlying layer), which is responsible for S. Unfortunately, there exist about as many accounts of "mechanism" as there are contributors to the debate. To some, a "mechanism" is the embodiment of the causal powers of a structure and thus something in the world (e.g., the critical realists), to others it is a piece of theory (see, *e.g.* Stinchcombe 1993 [1991]). Some are explicit methodological individualists and thus demand that mechanistic explanations cite (only) the interaction of individuals and their constraints (*e.g.* Little 1998) while others are explicit methodological holists (*e.g.*, Lawson 1997). As much as the details differ, "mechanism" always refers to that structure or process at the underlying layer which is causally responsible for the event or phenomenon at the empirical level (or, in some accounts, to a representation of the underlying structure or process). Synoptically, we can say that NMP subscribes to the following theses:

1. Theoretical explanation is the only or ultimate aim of social science.
2. The explananda of interest are empirical phenomena.
3. The desired explanantia consist of or make reference to the underlying causal mechanism(s) that give rise to the empirical phenomena.

Consider the following examples of causal-mechanistic explanations. The empirical phenomenon in our first case is that of bank runs. Robert Merton uses a well-known mechanism pattern in his explanation, the "self-fulfilling prophecy" (Merton 1968 [1948]). In general terms, a self-fulfilling prophecy is a mechanism by which an expectation induces just that kind of behaviour that is likely to make the expectation come out true. In this case some depositors (possibly, falsely) expect a bank insolvency and withdraw their savings. Observing the first withdrawers, other depositors start sharing the negative expectation, follow them and also withdraw their money. Thus expectation and withdrawing behaviour are mutually re-enforcing. Eventually, the bank can go bankrupt even if it was financially sound when the rumour started.

In the second case the empirical phenomenon is the well-established strong correlation between monetary growth and nominal income growth. Milton Friedman and Anna Schwartz try to explain this phenomenon by means of the following mechanism (Friedman and Schwartz 1963, pp. 60ff.). The initial change in the rate of growth of the money stock (produced by, say, a change in the rate of open-market purchases by the Fed) will induce investors to have more liquidity in their portfolios than they prefer. They will therefore seek alternative investments, initially in assets that are similar to the ones sold to the Fed, *i.e.*, relatively risk-free assets. The increased demand for low-risk assets will bid up their prices, and hence make higher-risk investments relatively more attractive. The initial impulse is thus spread out to various asset categories, including non-financial assets. The general increase in the prices of assets means an increase in wealth relative to income, which makes the purchase of goods and services cheaper (relative to the purchase of the sources of those goods and services, *i.e.*, assets). In this way, demand for goods and services will be increased, and therefore money income. With this story, of course, Friedman and Schwartz want to argue the monetarist case, *viz.*, that the causal direction (at the aggregate level) runs from money to nominal income and prices rather than the other way around (or mainly in this direction).

The third case concerns the relatively poor performance of the UK productivity growth rate in the past one hundred years (Lawson 1997, pp. 255ff.). Here the relevant underlying structure is the UK collective bargaining system, which is far more highly decentralised than in other countries. This structure "facilitates mechanisms which inhibit coordinated decision-making and works against quick or smooth responses to changes in production possibilities" (p. 257). The main difficulty consists in the higher degree of skill specialisation resulting from the decentralised system. Any change in technology, which makes a certain set of skills redundant, will thus meet a relatively stronger degree of resistance than in economies with a lower degree of specialisation. But this means that new technologies take more time to be implemented and hence productivity will grow at a slower pace.

III

As pointed out above, practicing social scientists seem to pursue a variety of different goals. Apart from theoretical explanation, I mentioned description, prediction and control. If my statement of NMP is correct, why can it ignore these alternative aims? There appear to be four possible reasons. A methodologist can ignore an aim X if (a) X is not important; (b) X is not attainable; (c) X is of no particular methodological interest because how to realise it is already well understood; (d) although he focuses on a different aim Y, he recommends a practice that helps realising both X and Y simultaneously.

This section considers the first three arguments and the next section the fourth. I will argue that none of the four arguments is successful for the non-explanatory aims of description, prediction and control. Let us consider each argument in turn.

Description, prediction and control are not important. One might argue that a given goal is not important either because no one in fact pursues it or because one ought not pursue it. As pointed out

above, as a matter of fact social scientists do pursue a number of non-explanatory goals. Thus, from a purely descriptive perspective it is simply false to deny description, prediction and control significance. But maybe this is the wrong approach altogether: one may argue that social scientists uniquely or ultimately *ought* to aim at explaining social phenomena. If practice falls short of this ideal, this only shows that there is something wrong with the social sciences, not with methodology.

Where do normative aims come from? Surprisingly, there is very little indication in the NMP literature about why (perhaps, ultimately) we should take only theoretical explanation seriously. As far as I can see, most proponents simply *presuppose* that the only aim worth pursuing is that of theoretical explanation. Perhaps this is a reaction to a number of anti-realist positions in the philosophy of social science, positions which tended to de-emphasise explanation as a significant goal.

It is clear, however, that methodologists should not prescribe scientists which goals and values to hold. Apart from prescriptivism and pure descriptivism, there is a third position that is normative, and yet does not impose philosophers' values on scientists. According to this third position, the practice of science should reflect the goals and values of those people affected by it. That is, as Kitcher 2001 argues, it should ultimately be up to the population to decide which research questions are worth pursuing, with what methods and how results are to be implemented. The aims of science, then, ultimately require democratic legitimisation. Kitcher calls a science that reflects the goals and values of people in its lines of research, its methods and dissemination "well-ordered science".

It is clear that current science is very remote from the ideal of full democratic legitimisation, and Kitcher himself admits that the specific model he sketches might not be practicable. But a few considerations show that it would be foolish to deny that description, prediction and control are likely to be part of "well-ordered" social sciences.

Accurately describing aspects of socio-economic systems is often regarded as a preliminary step to the worthy kinds of scientific activities. While it is certainly true that description sometimes plays a preparatory role, it is also an important end in itself. Especially when social justice is concerned, minimal differences in measurement procedures can lead to great changes in policy results. Consider the two concepts "inflation rate" and "rate of unemployment". The first measures the degree to which the general level of prices—usually but not always thought of as consumer prices—increases over a period of time; the second measures the proportion of unemployed to total members of the workforce. Changes in the measurement procedures associated with both concepts can cause great differences in policy results without concomitant changes in either the policy pursued or our understanding of the causal connections among the relevant variables because the concepts occupy a key place in established policy procedures. One reason for why policy makers need accurate estimates of the rate of inflation is that a great number of contracts and payments are inflation indexed. It is sometimes the intention of policy makers to keep the purchasing power of payments such as pension payments constant through time. In other words, the intention is to increase periodical payments by the amount of inflation for that period. The attempt to achieve this objective is frustrated if the price index does not accurately reflect the actual change in the general level of prices. Robert Gordon makes this point forcefully (Gordon 1992, p. 42f.):

> Price measurement matters. This is particularly true for errors in the rate of inflation over periods of three months to two years. In the American case, the late 1970s and early 1980s provided a textbook case of the follies that occur when policymakers rely on flawed price indexes. [...] Not only were millions thrown out of work in 1980 and again in 1982, but the distribution of income was permanently distorted as millions of others receiving social security benefits escalated by the CPI enjoyed a windfall gain, since their indexed payments went up by more than the true inflation rate. The unnecessary extra benefit payments caused the government deficit to increase, and higher indexed union wages put more pressure on firms to raise prices and continue the inflation.

Similarly, changes in the definition of the rate of unemployment can have automatic policy consequences in welfare states. If, for instance, long-term unemployed or people over sixty or

under eighteen are excluded from the definition of unemployment after a revision of the counting method (as happened in the UK during the 1980s), they might either lose benefits altogether or receive a different kind of benefit. In such a case, therefore, a change aimed primarily at improving the accuracy of an index can have real consequences for the people concerned.

Success and failure of our attempts to predict and to control social phenomena, too, are of obvious significance for many people. Any decision (by an individual, a company or a community) about an action whose outcome depends on the future value of some variable, would enormously profit from the ability to accurately predict that variable. I'd be rich if I could predict next year's stock market; the entrepreneur could plan his production very efficiently if he could predict future demand and input prices; the community could build sufficiently many roads, hospitals, kindergartens and power plants if it could only predict future traffic, number of inpatients, fertility rates and energy demand. For policy making we would further like to have knowledge that allows us to *control* social phenomena. Think of the "golden triangle" of economic policy—full employment, low or zero inflation and a sustainable rate of growth. Generations of politicians have spent their careers trying to achieve this triple goal, with of course a very mixed success record. Or think of more socio-political goals such as the elimination of poverty, the control of crime, the co-maximisation of liberty and security and so on.

From the point of view of policy making, description, prediction and control are enormously important goals of the social sciences. Importantly, they are scientific goals in themselves, not instrumental for a more ultimate scientific goals such as explanatory adequacy. (Of course, most ends are mere means from the point of view of another end: descriptive accuracy may serve social justice, which in turn may serve happiness and so on. My point here is that there is no immediate other scientific end such as explanation or understanding for which description, prediction and control are instrumental.)

Prediction and control are unattainable. Few people would deny that predictive success and success at controlling social phenomena would be enormously useful for practical and political purposes—if we knew how to reach them. The argument one hears frequently is a rather different one: we should not aim at predicting and controlling social phenomena because doing so with success is impossible. Immanuel Kant argued that "ought" implies "can"; if he was right, then "cannot" implies "ought not". (I am not aware of anybody who denies that social phenomena can be accurately described, hence I omit description in the present discussion. I will consider it below, in the context of methodology.)

It is fairly obvious, however, that the across-the-board assertion that prediction of social phenomena can never be successful is plainly mistaken. Pessimists tend to take predictions made in highly regimented physical systems as benchmark for predictions in "natural" (*i.e.*, unregimented) social systems. If that is the standard, of course the endeavours of social scientists must fail. The comparison is obviously unfair. It is neither the case that prediction in natural physical systems always works (just think of weather and other complex or chaotic systems) nor that prediction in artificial social systems doesn't (think of some of the results of experimental economics). Further, the fact that predictions in social science fare relatively poorly (relative to what standard?) should not prevent us from trying as long as we can improve upon mere guesswork.

And that we can improve on guesswork is impressively demonstrated by recent work in financial econometrics. Out of the many examples one could cite, I find this one particularly telling because the belief that financial time series are unpredictable is widespread. Common wisdom holds that if financial time series were predictable, someone would predict them because he could profit from it. His action, however, would destroy the basis for the predictability of the series for others. Since this is true for everyone, nobody can systematically predict financial time series. This story underlies the old economists' joke, which is (Lo and MacKinlay 1999, p. 6)

> about an economist strolling down the street with a companion when they come upon a $100 bill lying on the ground. As the companion reaches down to pick it up, the economists says "Don't bother—if it were a real $100 bill, someone would have already picked it up."

For financial markets, this means that price changes must be unpredictable, *i.e.*, that financial time series are random walks. Paul Samuelson first applied the Random Walk Hypothesis to financial markets in a theoretical paper (Samuelson 1965). Samuelson's and others' theoretical work (in particular, Fama 1970) has led to the belief that "a blindfolded monkey throwing darts at a newspaper's financial pages could select a portfolio that would do just as well as one carefully selected by the experts" (Malkiel 1991, p. 24).

Since the original formulation of the Random Walk Hypothesis, literally thousands of articles with countless theoretical and empirical results have appeared. Importantly, many of these studies are inconsistent with the hypothesis. In particular, it could be rejected to hold for some markets (see, *e.g.*, Lo and MacKinlay 1988 on U.S. securities), which suggests that carefully designed active management strategies can outperform mere guesswork in the long term. Of course, this does not mean that making excess profits in these markets is easy. But it does show that investing in forecasting methodology can sometimes pay off.

Whereas prediction is a two-place relationship between a set of information and a set of target variables, control is a three-place relationship between a set of information, a set of actions and a set of target variables. When we say we can control a variable, we mean that our actions are able to bring about a specific value of that variable, a value that it would not have taken in the ordinary course of things. Control is thus prediction in a specific context: we predict the value of the target variable not upon passively observing the information set but rather upon an action designed specifically such as to bring about that value.

Thus, if successful prediction is already impossible, successful control will be impossible *a fortiori*. Consequently, arguments to the effect that all efforts to conduct socio-economic policy are futile abound. Especially following the Lucas critique, many models in the rational expectations literature were used to demonstrate the ineffectiveness of economic policy.

Again, it is crucial here to get the standard right. Control in open, largely unregimented and living systems will be harder than in static systems that can be totally shielded from outside influences and where no ethical considerations matter. However, there are many examples of successful policies—just think of the effective control of inflation worldwide since the 1970s, the success of the crime-control programme in New York City in the Giuliani era or the skill with which Alan Greenspan has been steering the U.S. economy.

Consequently, none of the principled arguments that social technology can never work is convincing. Laissez-faire policies are sometimes justified on the grounds of some of Friedrich von Hayek's work. But all Hayek really said was that across-the-board social planning would have to ignore vast amounts of information that markets generate and thus likely lead to inferior results. He wasn't an opponent of piecemeal social engineering (though he did not like this Popperian term, see, *e.g.*, Hayek 1982, vol. III, p. 204). Similarly, the Lucas-critique is sometimes taken to demonstrate that economic policy based on econometric results must be ineffective. However, the Lucas critique is based on the assumption that agents have information about all causally relevant variables at their disposal, including information about unpredictable future events. This is a preposterous requirement. In a more realistic setting, agents will chose the best forecasting model among a number of alternative models, all of which will be mis-specified in one way or another. Now, as we shall see in more detail below, if models are mis-specified and if underlying structures frequently change, simpler, more adaptive models can be more appropriate forecasting tools than causally more adequate models. Sensible agents will use robust forecasting rules. However, a world where agents use simple, adaptive forecasting devices is immune to the Lucas critique (see Hendry 2002).

Description, prediction and control are methodologically trivial. The third argument practically claims the opposite of the second: these aims are not impossible to reach but rather uninteresting because

how to reach them is already well understood. In fact, however, all three areas are methodological minefields.

In the context of description, especially economics suffers from a well-recognised but by no means resolved problem: the variables constructed by statistical offices tend to differ quite significantly from the variables that figure in economic theories. This makes the application of theories to data—be it for explanation or prediction or policy—problematic to say the least (see, *e.g.*, Stigum 2003 for a discussion of the problem).

Suppose a theory predicts that some variable X causes another variable Y and we want to test that theory by running a regression of Y on X. For the moment suppose away all problems about causal confounders. Still, the regression coefficient of X will be biased if X is measured with error. The standard response is to find a so-called instrumental variable Z such that (a) Z is correlated with X but (b) uncorrelated with the measurement error. Including Z in the regression allows to measure the coefficient of X without bias. But: it is not guaranteed and, in fact, quite unlikely that a suitable instrument can be found. The error term is by its very nature unobservable. How, then, could we know that some other variable is uncorrelated with it? Consequently, measurement error is widely regarded as a severe problem for empirical modelling. And yet, it has attracted very little interest from methodologists.

Even if the problem is not in the relation between empirical and theoretical variables, accurate measurement is methodologically intricate. One case I've been working on in some detail concerns the measurement of inflation. Inflation is the rate of change in an index which measures the aggregate level of prices in a country or region. The so-called index number problem is the problem of how to aggregate a large set of individual prices (and, perhaps, traded quantities) such that a single number results. Essentially, thus, index numbers are averages, and the associated problem is to decide which of an infinity of possible averages represents the quantity "price level" most accurately.

In the early to mid-1990s rumours accumulated that the U.S. Consumer Price Index CPI overstates "true" consumer price inflation about 1 percentage point annually. As mentioned above, changes in the definition of inflation can have dramatic policy consequences because many spending programmes as well as tax brackets and private contracts are inflation indexed. It was estimated at the time, for instance, if the CPI did indeed overstate consumer price inflation by about 1 percentage point, this bias would contribute $146 billion to the budget deficit in 2006 and $691 billion to national debt by then. Consequently, the Senate Finance Committee established a panel of experts to study the accuracy of the CPI. This commission, aka the Boskin commission (after its chair, Stanford economist Michael Boskin) published its report at the end of 1996 and estimated a bias of, indeed, 1.1% per year.

As can be expected, the report did not only receive acclaim. One response I found revealing was a short book by the economist Dean Baker who re-estimated the bias on the basis of a similar methodology but a slightly different set of assumptions (Baker 1997). Example: the Boskin commission observed that discounters gain market share and interpreted this as an indication that consumers regard the reduction in services and convenience as more than offset by the lower prices. Consequently, the commission believed the CPI to be biased to the extent that discounters are underrepresented in its outlet mix. Baker interprets the same phenomenon as an indication that consumers are forced to shop at discounters due to increased poverty, estimates that the quality losses are not compensated by reduced prices and concludes that the bias due to outlet substitution is negative rather than positive (*i.e.*, that the CPI understates rather than overstates inflation in this category).

What Baker's study, as well as many others, in my view show is that at least the following aspects of inflation measurement are far from being well understood:

- what is the right measurement objective for the CPI?
- how should one deal with product substitutions (*i.e.*, when consumers replace a good whose relative price has increased by one whose relative price has decreased)?

- how should one deal with outlet substitutions (*i.e.*, when consumers replace traditional distribution channels such as grocers with discounters or internet-based traders)
- how should one deal with quality changes (*i.e.*, when the quality of a good represented in the index changes)?
- how should one deal with new goods (*i.e.*, when goods appear in the market that don't have a direct predecessor such as mobile telephones)?
- how should one deal with taste changes (when for example consumers' preferences change due to changes in income)?
- how should one deal with the fact that different groups differ in salient respects (*e.g.*, employees, old people and poor people)?
- should expert knowledge play a significant role in price measurement, and if so, which?

Many replies to these questions will make a difference to the value of the CPI and hence to the disposable incomes of countless people and the federal budget. To repeat myself here: price measurement matters.

The work I know best in the area of forecasting and policy analysis methodology is that of David Hendry and his collaborators. Since Hendry *et al.* have published a number of survey articles about these topics, here I will merely summarise what they regard as the most pressing open methodological issues. In a survey of recent work on forecasting, Hendry and Michael Clements list the following ten areas, the increased understanding of which is crucial for our forecasting success (Hendry and Clements 2003, p. 315):

- Pre-testing for the inclusion of intercept corrections
- Modelling shifts
- Inter-forecast smoothing
- The role of survey information in forecasting
- Pooling of forecasts and forecast encompassing
- Discriminating measurement errors from innovation shifts
- Multi-step estimation for multi-step forecasting
- The advantages of explicitly checking co-breaking for forecasting
- Attempts to forecast rare events
- Leading indicators in forecasting.

As hinted at above, control is forecasting in a special scenario, *viz.*, when a policy variable is deliberately changed. Consequently, some of the methodological issues concerning forecasting are relevant for control, too. Particularly relevant is the role of co-breaking relationships. The time series of two variables are said to be co-breaking if both are subject to deterministic shifts (*e.g.* to shifts in their unconditional means) but a linear combination removes the shifts. When a policy variable and a target variable co-break, the link can be exploited to control the target variable (see Hendry 2004).

Although the tests for co-breaking relationships have been developed, they remain unsatisfactory because they are oversized, *i.e.*, the probability of erroneously concluding that no co-breaking is present is higher than the chosen significance level suggests, especially when the degree of autocorrelation is high. Further, all results are asymptotic, which indicates that they have unfavourable small-sample properties (Hendry and Massmann forthcoming). All that said, one can conclude that methodological problems in measurement (or description), prediction and control abound.

The fourth and perhaps strongest argument for why a methodologist might be justified in ignoring an aim of science is that the practice one recommends simultaneously achieves a plurality of aims. In the present context, one might argue that the focus on explanation is innocuous because investigating causal mechanisms will also further the other ends of the social sciences. In this section I consider whether investigating causal mechanisms is always a good strategy when the aim is accurate description, successful prediction or control.

Causal-mechanistic models are the best models of data. Let us return to the measurement of inflation for a moment. There are three approaches to the abovementioned problem of index numbers: the axiomatic, the stochastic and the economic approach. Of these three, only the last approach makes use of causal-mechanistic hypotheses by regarding the problem as one of expenditure minimisation given a vector of commodity prices. According to this theory, the right answer to the question "Which is the best index?" is "That which measures the change in expenditure for the basket of goods optimising agents would choose were they to maintain a constant level of utility." The other two approaches are based on very general mathematical and statistical considerations and can therefore be regarded as non-mechanistic.

In this context we could unify our interests in theoretical explanation and accurate description if we could demonstrate that the economic approach is the best answer to the index-number problem. There are, however, two problems with this strategy.

It is true that among economists the economic approach is widely accepted as providing the right answer to the index-number problem. However, and this is the first problem, conceptually it is not clear whether utility is the right quantity to be held constant in a price index. One of the purposes for which the CPI is used is to insulate benefits recipients from changes in the cost of living. But the "cost of living" has other connotations than the "cost for achieving a given level of utility". Another is the "cost for achieving a given set of functionings" (in Amartya Sen's sense, see Sen 1999). Yet another requires to fix the relative position in society. Yet another asks to define a basket of goods that yields "acceptable standard of living". It is far beyond the scope of this paper to attempt to resolve such normative issues in the foundations of economics. However, it is easy to see that claiming that only the economic approach provides the correct answer to the index-number problem begs the question against these alternative understandings. And not all of the alternatives will make use of mechanistic information in constructing an index number. A strict price index, for example, which traces the cost for purchasing a fixed basket of goods, is ignorant as to why consumers purchase this rather than another basket.

Second, even if we accept that fixing utility is the right answer to the index-number problem, we can face a trade-off between two different desiderata. For policy purposes, it matters greatly that the index is readily measurable. But a constant-utility index has one major deficiency: it is unobservable. An exact constant-utility index can be measured only if the preference functions of the consumers are known. In practice, of course, they are not. In order to solve this quandary, the effects of various changes that happen in the market are imputed by government statisticians or other experts. If a new good appears, for example, it is estimated whether it is an equivalent replacement of the old good or whether it provides a different amount of utility and how much (the price change is then discounted by the amount of the utility change before it enters the index). But there is something deeply dissatisfactory with this practice. One reason behind the indexation of federal spending programmes was to introduce a mechanical procedure where there had been fortuitous decisions of politicians. Demanding that the CPI be a constant-utility index re-introduces the earlier uncertainty as now statisticians and other experts have to decide how much the cost of living has changed. A constant-utility index may be theoretically more satisfactory but this advantage is bought at the expense of much greater measurement uncertainty. Angus Deaton, for example, argues (Deaton 1998, p. 38):

> In fact, it is unclear whether a quality-corrected cost-of-living index in a world with many heterogenous agents is an operational concept. I argue that the major problem is not the Consumer Price Index itself, but what it is used for; it is foolish to index benefits thoughtlessly and mechanically to a concept that is hard to define and harder to measure.

If this is true, there are good reasons to believe that causal-mechanist models aren't always the best models of data. True, they may be more satisfactory from the point of view of theoretical explanation. But the satisfaction of this desire is purchased at the cost of other desiderata, which in some circumstances may be as or more important than theoretical explanation.

Causal-mechanist models are the best models for prediction. This is a version of the old symmetry thesis according to which (causal) explanation and prediction are the two sides of the same medal. Since the symmetry thesis has long come out of fashion I don't really have to argue that causal-mechanist models will not always be the best models for prediction. It is useful, however, to consider some recent technical results in the theory of forecasting.

The older theory of forecasting was built on two presumptions: (a) that the econometric model provides an adequate representation of the underlying data-generating structure; (b) that that structure remains stable within the forecasting horizon. In more recent work in this area, it has come to be accepted that (a) econometric models are almost always mis-specified (*i.e.*, they do not represent the underlying structure correctly); (b) socio-economic systems are subject to frequent structural breaks. Under these conditions, it cannot be proved that causally more adequate models beat models that are not based on causal variables in forecasting competitions. The vital property is that of adaptability: a model that adapts more rapidly to a structural break will beat a model that, after the occurrence of a break, is permanently off track. But since (as of today) causal models tend not to be robust to such shifts, non-causal models often outperform them (Hendry and Mizon 2001, p. 5):

> A key consequence of these results is that the best available forecasting model need not be based on the 'causal determinants' of the actual economic process, and as the example… shows, may be based on 'non-causal' variables, that is, variables which do not enter the DGP [*i.e.*, the data-generating process or *mechanism*].

Causal-mechanistic models are the best models for policy. This idea is based on the current version of the symmetry thesis. In this version, the symmetry is between causal explanation and intervention or control. The idea is, roughly, X causes Y if and only if an intervention (of the right kind) that changes the value of X changes the value of Y (Woodward 2003). That is, according to this theory, if Y is a target variable of interest and X is a policy variable we can manipulate, then if and only if X causes Y, we can use X in order to control Y. Values of X cause values of Y and thus, on this account, X explains Y. At the same time X figures as a policy variable through which we can change Y.

The problem with this approach in the present context is that an intervention, in order to serve the appropriate role in a theory of causation and explanation, has to have a number of very idealised properties. The theory states that IF we were to intervene in just the right way and changed the value of the first variable, then the value of the second variable would change if and only of the first caused the second. This formulation is needed because we want to say that, say, past events or event-types we cannot manipulate in the right way are nonetheless causes of other events. However, for policy we need real interventions, not hypothetical ones. Hence, even if it is true that some X causes some Y, we might not be able to manipulate X in just the right way such that a change in Y ensues.

Consider James Woodward's definition of an "intervention" (Woodward 2003, pp. 98f.):

> I1. I causes X

> I2. I acts as a switch for all other variables that cause X [*i.e.*, I breaks all causal laws that have X as an effect]

> I3. Any directed path from I to Y goes through X [*i.e.*, if I causes Y, it does so only through X]

I4. *I* is (statistically) independent of any variable *Z* that causes *Y* and that is on a directed path that does not go through *X*

I5. *I* does not alter the relationship between *Y* and any of its causes *Z* that are not on any directed path (should such a path exist) from *X* to *Y*,

where a "directed path" is a sequence of direct causal relationships (*e.g.*, such that V_1 causes V_2 causes … causes V_n). Following this definition, let us call an intervention that satisfies these conditions a "Woodward-intervention". This contrasts starkly with ordinary usage, which is stronger and weaker at the same time. According to ordinary usage, it is central that (R1) *I* can be set at the will of a human agent (an experimenter, a policy maker or what have you) who aims at changing a policy variable *X* and (R1) *I* sometimes succeeds in changing the value of *X*. If, say, a central bank intervenes on short term interest rates (*X*) by conducting open-market operations (*I*), it can, on occasion, fail to influence *X*. However, we would not call a variable an intervention on *X* if it never affected *X*. Praying that a dear friend recovers from an illness is not an intervention. Let us call an intervention that satisfies R1 and R2 a "real intervention".

With the possible exception of I1 (depending on how we interpret the word "cause" in I1), any of Woodward's conditions can be violated by a real intervention, and some conditions are very likely to be violated. Few interventions will cancel all causal influence other variables have on the policy variable. An open-market operation is likely to *influence* the short-term interest rate but it won't remove all of its other causes. Some interventions will affect the target via more than one route. Suppose smoking is a policy variable and mortality rates due to lung cancer the target. Banning smoking in bars might affect smoking. But it might, too, change the frequency with which people go to bars, thus trigger alternative activities which may in turn influence mortality rates due to lung cancer. Some interventions will be correlated with other causes of the target variable. The level of the minimum wage in a state is sometimes adjusted in response to favourable economic conditions. If salary is regarded as the policy variable and employment as the target, here the intervention is correlated with another cause of the target, *viz.*, "economic conditions". Further, many interventions will affect the causal laws of the system themselves. The oil price shocks were interventions that have changed many consumers' and producers' attitudes towards petroleum use and are thus likely to have altered the economy's causal laws in important respects.

Therefore, even if *X* causes *Y* in Woodward's sense, the relationship will not always be stable under a real intervention on *X*. What matters for policy is the stable association between the policy variable and the target, not the reason *why* the correlation is stable. Suppose we believe *X* to cause *Y* and *I* to be an intervention, both in Woodward's sense. But let the real structure be such that the correlation is due to a common cause *Z*, and let *I* affect *Z* or, alternatively, *X* and *Y* independently, perhaps such that *X* changes before *Y* changes, so that it looks as if *X* causes *Y*. There is nothing wrong with this from a policy point of view. But it would be mistaken to explain the change in *Y* by citing the change of *X*. There is hence a third notion of intervention, that of a policy-intervention. A "policy intervention" adds to the "real intervention" the requirement that the relationship between policy variable and target remain invariant (*cf.* the econometric notion of super exogeneity; see, *e.g.*, Engle *et al.* 1983).

Furthermore, the locution "there is a causal mechanism between *X* and *Y*" (in the sense relevant for this paper) is considerably weaker than the locution "*X* causes *Y*" (in Woodward's sense). An intervention (real or Woodward) may, for example, trigger two different mechanisms whose influence on the target variable can be positive, negative or nil. Consider following example from economist/social theorist Julian LeGrand (LeGrand 1982, p.148, quoted from Elster 1999, p. 50):

> A high marginal tax rate lowers the opportunity cost or "price" of leisure, and, as with any commodity whose price is reduced, thereby encourages people to consume more of it (and thus do less work). But, on the other hand, it also lowers peoples' incomes, and thereby may induce them to work harder so as to maintain their standard of living. These two effects—the substitution and income effects, in economists parlance—operate in opposite directions, and their net effect is impossible to predict from theory alone.

After the fact, we'll be able to explain the change (if there is any) by citing the two mechanisms. But (in this example) the marginal tax rate does not appear to be a good policy variable if employment is the target.

Another reason for why it will not always be enough to have a mechanism that connects X and Y for using X as a policy variable is that causation in Woodward's sense is not always transitive: if X causes Z and Z causes Y, it is not guaranteed that X causes Y. In the social sciences, threshold effects can illustrate possible failures of transitivity. A threshold effect obtains when a cause changes the value of its effect only after reaching a certain critical value. This phenomenon has been reported in the relation between human capital and income, for instance. Human capital needs to reach a critical level in order to affect income. That is, there is a Woodward-intervention on human capital that changes the value of income, *viz.*, one that raises human capital from below to above the threshold level (and thus human capital causes income). Further, let there be a Woodward-intervention on schooling that changes the level of human capital (and thus schooling causes human capital). However, it is well possible that no intervention on schooling affects human capital enough in order to get it above the threshold value (and thus, schooling does not cause income). Hence, we cannot use schooling as a policy variable to control income as the target—despite the causal mechanism running from the policy to the target variable.

In summary, mechanisms and control are distinct albeit related. Both are causal notions to be sure. But they are causal notions with different characteristics and hence investigating whether one applies will not always be the best strategy for finding out whether the other applies.

V

For reasons discussed above, David Hendry and his collaborators emphasise the distinctness of forecasting and policy (Hendry and Clements 2003, p. 314): "Thus, neither forecast success nor failure entails either good or bad policy advice: policy models need policy evaluation". It seems appropriate to add: data models need data evaluation; forecasting models need forecasting evaluation; and, of course, explanatory models need evaluation in terms of explanation. The social sciences come with a variety of different ends, and each end brings its own standard of evaluating models.

Maybe I am preaching to the choir here because we've all become pluralists already. Perhaps. But if that is so, it should strike one as odd that measurement, forecasting and policy evaluation are virtually absent from the debates in methodology and philosophy of the social sciences—and that despite the importance and methodological intricacy of these activities. And there is at least one significant movement in the philosophy of the social sciences that is quite explicit in its neglect or rejection of non-explanatory aims. So maybe some methodologists' declared belief in pluralism is little more than lip service.

Against critical realists and other proponents of the "new mechanist perspective" I therefore want to urge that methodology and the philosophy of the social sciences can profit from a more explicit endorsement of various non-explanatory aims of the social sciences. Of course, far be it from me to make the same mistake and de-emphasise the explanatory role of social theory. But, to stress it again, explanation is only one important endeavour among many.

Bibliography

Baker, Dean (ed.) 1997, Getting Prices Right: The Battle Over the Consumer Price Index, Armonk (NY): M.E. Sharpe.

Brante, Thomas 2001, "Consequences of Realism for Sociological Theory-Building", Journal for the Theory of Social Behaviour 31 (2): 167-194

Deaton, Angus 1998, "Getting Prices Right: What Should Be Done?" Journal of Economic Perspectives 12 (1): 37-46

Elster, Jon 1983, Explaining Technical Change: A Case Study in the Philosophy of Science, Cambridge: Cambridge University Press

Elster, Jon 1985, Making Sense of Marx, Cambridge: Cambridge University Press

Elster, Jon 1989, Nuts and Bolts for the Social Sciences, Cambridge: Cambridge University Press

Elster, Jon 1999, "A Plea for Mechanisms", in Peter Hedstroem and Richard Swedberg, (eds.), Social Mechanisms: An Analytical Approach to Social Theory, Cambridge University Press, Cambridge, pp. 44-

Engle, Robert, David Hendry and Jean-François Richard 1983, "Exogeneity", Econometrica 277-304

Fama, Eugene 1970, "Efficient Capital Markets: A Review of Theory and Empirical Work", Journal of Finance 25: 383-417

Friedman, Milton and Anna Schwartz 1963, "Money and Business Cycles", Review of Economics and Statistics 45 (1, Part 2, Supplement): 32-64

Hacking, Ian 1983, Representing and Intervening, Cambridge: Cambridge University Press

Hayek, Friedrich von 1982, Law, Legislation and Liberty, London: Routledge

Hedstroem, Peter and Richard Swedberg (eds) 1999a, Social Mechanisms: An Analytical Approach to Social Theory, Cambridge: CUP.

Hedstroem, Peter and Richard Swedberg 1999b, "Social Mechanisms: An Introductory Essay", in Peter Hedstroem and Richard Swedberg, (eds.), Social Mechanisms: An Analytical Approach to Social Theory, CUP, Cambridge

Hendry, David 2002, "Forecast Failure, Expectations Formation, and the Lucas Critique of Work", Economics Group: Economics Papers, 2002-W8, Nuffield College, University of Oxford

Hendry, David 2004, "Causality and Exogeneity in Non-Stationary Economic Time Series of Work", Causality: Metaphysics and Methods Technical Reports, CTR 18/04, CPNSS

Hendry, David and Michael Clements 2003, "Economic Forecasting: Some Lessons From Recent Research", Economic Modelling 20: 301-29

Hendry, David and Michael Massmann forthcoming, "Co-Breaking: Recent Advances and a Synopsis of the Literature", Journal of Business and Economic Statistics

Hendry, David and Grayham Mizon 2001, "Forecasting in the Presence of Structural Breaks and Policy Regime Shifts of Work", Economics Group: Economics Papers, 2002-W12, Nuffield College, University of Oxford

Kitcher, Philip 2001, Science, Truth and Democracy, Oxford: OUP

Lawson, Tony 1997, Economics and Reality, London: Routledge

Lawson, Tony 2003, Reorienting Economics, London: Routledge

LeGrand, Julian 1982, The Strategy of Equality, London: Allan and Unwin

Little, Daniel 1991, Varieties of Social Explanation, Boulder (KA): Westview

Little, Daniel 1998, Microfoundations, Method, and Causation: On the Philosophy of the Social Sciences, New Brunswick (NJ): Transaction Publishers

Lo, Andrew and Craig MacKinlay 1988, "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test", Review of Financial Studies 1 (41-66)

Lo, Andrew and Craig MacKinlay 1999, A Non-Random Walk Down Wall Street, Princeton: Princeton University Press

Malkiel, Burton 1991, A Random Walk Down Wall Street, 5th ed. ed. New York (NY): Norton

Merton, Robert 1968 [1948], "The Self-fulfilling Prophecy", in Robert Merton, (ed.), Social Theory and Social Structure, Free Press, New York, pp. 475-90

Samuelson, Paul 1965, "Proof That Properly Anticipated Prices Fluctuate Randomly", Industrial Management Review 6: 41-9

Sayer, Andrew 2000, Realism and Social Science, London: Sage

Sen, Amartya 1999, Development As Freedom, Oxford: OUP

Stigum, Bernt 2003, Econometrics and the Philosophy of Economics: Theory-Data Confrontations in Economics, Princeton: Princeton University Press

Stinchcombe, Arthur L. 1993 [1991], "The Conditions of Fruitfulness of Theorizing About Mechanisms", in A. Sørensen and S. Spilerman, (eds.), Social Theory and Social Policy: Essays in Honor of James S. Coleman, Praeger, Westport (CT), pp. 23-41

Woodward, James 2003, Making Things Happen, Oxford: Oxford University Press