

CLARK GLYMOUR

A SEMANTICS AND METHODOLOGY FOR *CETERIS PARIBUS*
HYPOTHESES

ABSTRACT. Taking seriously the arguments of Earman, Roberts and Smith that *ceteris paribus* laws have no semantics and cannot be tested, I suggest that *ceteris paribus* claims have a kind of formal pragmatics, and that at least some of them can be verified or refuted in the limit.

1.

Many, perhaps most, of the results of scientific inquiry are generalizations towards which we have a peculiar attitude: although their apparent logical form is general, and we assert them and use them in prediction, in designing experiments, in deriving new relationships, and in rejecting contrary hypotheses, they have actual or possible (feasibly possible) counterexamples, we know they have counterexamples, and we cannot say precisely under what conditions counterexamples will not arise in the future. The claims hold in “normal” conditions, whatever those are, but not otherwise. Earman, Roberts and Smith (*this issue*; hereafter ERS) call such claims “putative *ceteris paribus* laws”. Putative *ceteris paribus* laws occur, for example, in almost everything we know about cellular biology, in all of the causal claims of the social sciences, and throughout the medical sciences, ERS say three things about such claims: they are not laws, we do not understand what they say (they have no “semantics”) and they cannot be tested. So far as I understand ERS, their arguments are: *ceteris paribus* generalizations are not laws because they are vaguely qualified by “normally;” they have no semantics because we have no general specification of the meaning of “normal”, and they cannot be tested because there is no description of a possible detectible event, or sequence of events, with which they are inconsistent – counterexamples are, *ipso facto*, not normal.

Partially in response to these challenges, also in this issue, Spohn (*this volume*) offers a brilliant formal analysis of laws as pairings of general claims and ordinal rankings of propositions, rankings that represent dispositions, the disposition not to surrender the predictions of the generalization in the face of counterexamples. To test such laws, or at least to learn about



them, he suggests, is just to alter, by an appropriate procedure, the strength of one's dispositions to persist in trusting the predictions of a generalization as examples and counterexamples accrue. Along the way he gives an elegant, if perhaps unsatisfying, explanation of the "normal" conditions for a *ceteris paribus* law: the disjunction of all of the conditions in which the generalization holds.

I leave it to ERS whether Spohn's account responds to their arguments. I have a different take on the issues about belief and learning involved. I defer to Humpty Dumpty on whether *Ceteris paribus* generalizations, even those we believe, should be called "laws", but I think ERS do implicitly have a view of the pragmatics, if not the semantics, of such generalizations, and do implicitly have a view of testing or learning that can be made explicit. Whether it leads to their conclusions is another matter. In any case I shall foist it on them. When ERS say *ceteris paribus* claims have no semantics, I take them to mean that there is no intersubjective, general account of the sense of "normal"; that allows, however, for subjective accounts, for treating "normal" as a kind of indexical, and allows, if not a semantics, a pragmatics. When ERS say that *ceteris paribus* claims cannot be tested, I take them to be saying that there are no possible finite sequence of observations that decisively refute or verify such claims; that allows, however, that such claims might be refuted or verified in the limit. I propose to investigate the howevers.

2.

I take the point of view that the aim of inquiry is first cognitive, and only derivatively practical. The aim of inquiry is to come to believe interesting truths and to avoid believing falsehoods (practical matters, the necessities of action, influence what is of interest, of course); the business of epistemology is to say what attitudes we should have towards hypotheses generated in the course of such inquiry; and the purpose of methodology is to provide the most efficient, feasible and reliable procedures to fulfill the aims of inquiry. Spohn takes the second of these as his starting point; I take the third. They interact, of course. Testing, a part of methodology, is the business of specifying procedures that will yield detectible outcomes that will, or should if we follow a related procedure, alter our attitude towards hypotheses and take us a step towards the aim of inquiry.

There is a philosophical framework, originally due to Hilary Putnam but nowadays most developed in computer science, for proving things about how methods can and cannot lead to the truth, reliably, efficiently and feasibly. It is quite abstract (none of the juicy stories here about how

Darwin discovered evolution or about Faraday and lines of force) but has the virtue that it can be applied in almost any context.

The set up is this: There is a set \mathbf{H} of hypotheses, which may be infinite and which may not each be describable at some moment, but each of which can eventually be generated. Logically possible hypotheses not in \mathbf{H} constitute “background knowledge” – what is known not to be true, or assumed not to be true in the context of inquiry. For each hypothesis h in \mathbf{H} there is a set of possible data streams consistent with h . Each data stream is a sequence of objects, which may be descriptions or facts or events or situations or bangs on the sensory organs, but which are intended to represent the data obtained from experiment or observation. We allow that the same object may occur more than once in a data stream. I will assume that the data streams are unending, although with some technical reformulations that is inessential. The idea is that each element of a data stream is something that, if it occurs, an investigator following a method of inquiry detects. I will refer to a set of hypotheses and their associated data streams as a *discovery problem*.

In this setting a method of inquiry is simply a (partial) function whose domain is all initial segments of all data streams for all hypotheses in \mathbf{H} and whose range is \mathbf{H} . If the value of a method for a particular initial segment of data is h , I will say for brevity that the method *conjectures* h on that initial segment. The reliability, efficiency and feasibility of a method can have many forms. For feasibility, one can require that the partial function have any particular computational complexity one wishes; for efficiency one can require that a method get to the truth as fast as any method can in general, where speed can be measured, for example, by how often the method changes the hypothesis it conjectures.

My concern will be with reliability, which I will take to be this: A method is reliable for a discovery problem if and only if for every h in \mathbf{H} , and for every data stream d , the value of the method (remember, a method is a function from initial segments of data streams to \mathbf{H}) for all but a finite number of initial segments of d is h if and only if d is a data stream for h .

Why not require something stronger for reliability, for example that the method, which is a partial function, be correct on the first conjecture it makes – in other words, that the method remain silent until it has enough data to be sure which hypothesis in \mathbf{H} is true, and then report that hypothesis, and always be right? The answer is that for general, universally quantified claims, no such method exists that decides reliably between the claim and its negation. The point is in *The Meno*, and the proof is in Sextus Empiricus. The sort of reliability I defined first, usually called limiting

reliability in the literature, is the best we can do when we are talking, as we are here, of universally quantified claims.

The limiting reliability criterion just stated is equivalent, for most problems, to Bayesian convergence to the truth, provided no computational constraints are placed on the method. That is, there exists a reliable method for a discovery problem if and only if there are a prior probability assignment to the hypotheses and initial segments of data streams such that, conditioning on the data, the probability distribution converges to 1 for the correct hypothesis. In some deterministic discovery problems this is a great advantage, because one can describe procedures for conjecturing hypotheses that accord with Bayesian convergence criteria without the technicalities of distribution theory or worries about whether the priors have been specified appropriately, all of which turn out to be gratuitous. When, more realistically, computational constraints are imposed on methods, the equivalence disappears and the learning theory framework dominates. Even if one requires only that methods be Turing computable: there are discovery problems that are solved reliably, as described above, by Turing computable methods, but for which convergence to the truth in all cases is not possible with Turing computable probability measures. The converse is not true.

Learning theory is part of methodology, and it says nothing about what attitude one should have towards the conjectures of a method, even a reliable method. Most of epistemology – Bayesian epistemology, Spohn's ranking theory, sense datum theory, classical skepticism, whatever – is, in one way or another, about that question. There are connections, of course. Most skeptical arguments, for example, turn on elementary considerations about the nonexistence of methods to reliably solve formally simple discovery problems, and ERS present us with the skeleton of a skeptical argument about *ceteris paribus* generalizations. In the other direction, reliability demands can be imposed on schemes for changing ranking functions (just as on schemes for changing probabilities), and Kelly (1998), has shown that only a few (one of which is due to Spohn) of the many proposals meet that requirement. And there are many variations on learning theory that more closely entangle epistemology and methodology. Instead of giving particular hypotheses as output, we could consider methods that give probability distributions over hypotheses (posterior probabilities are just one possibility), or partial orderings over hypotheses (Spohn's orderings are an example), or any of a variety of scores that are used in practical statistics when posterior probabilities are infeasible to calculate, or scores that are a function of the number of data points since the method changed its value, etc. We can equally consider the advantages and disadvantages

of communities of inquirers using different methods for the same problem. That has been done for some while by computer scientists, by a few economists, and in philosophy, only by Philip Kitcher (1990).

Whatever the output of a method, the question arises as to its normative role in practical action, and only the Bayesians have a full, if highly idealized, story to tell: take the action with highest expected utility. (The recommendation is chiefly a philosophers' and economists' toy. I have never seen it done in any real scientific application, because even with explicit utilities it is computationally intractable in realistic cases unless one is very dogmatic about the possible hypotheses.) Endless alternatives are possible for forming propositional attitudes or making decisions using the output of a learning theoretic framework in which a single hypothesis is conjectured at a time, but I will advocate none of them here.

3.

Formal linguistics is not just semantics. It includes pragmatics as well, the analysis of speech acts whose truth conditions depend on the speaker, and more generally the analysis of indexical expressions whose denotations depend on time and place and context: *I, you, here, now*. Our understanding of language is not poorer, but richer, for recognizing and analyzing indexicals, even if we step out of formal "semantics" to something else.

Indexicals are possible in methodology as well. They are championed, in other terms, by philosophers who are *relativists*, who think that the truth value of a sentence in a language depends on something else besides the world and the language, on the endorsement of the community, or the conceptual scheme, or the method of inquiry, or something, and that in virtue of that something the truth values of sentences change in the course of inquiry. Sometimes this change is described as a change of language: the meaning of the sentence changes, depending on the something, while the syntax is kept, but I think the former way is easier to keep in mind without confusion.

Developments of relativist theories are usually placed within historical or sociological discussions, and their advocates are generally thought of as anti-formal. That is a pity, because formulated within learning theory they have an intricate structure, of interest to us here because it is a useful model for understanding *ceteris paribus* generalizations. Suppose we have a discovery problem that is ill-posed in the following respect: the set of data streams consistent with a hypothesis is itself a function of the conjectures of a method, or, put another way, for a given data stream, the hypothesis it corresponds to may change at any point in response to the conjecture

produced by a method. The formal details are developed in Kelly and Glymour (1992) and in Kelly (1996), and Kelly et al. (1992), but the idea should be clear enough.

The notion of limiting reliability of a method for a discovery problem still makes sense when relativized, in fact several distinct senses. In one sense of reliability, a method is reliable if there comes a time after which truth values cease to jump around and there also comes a time after which the method is always correct in its conjectures under the stable truth values. More interestingly, the truth values may change without end, but a reliable method eventually changes in twain and after a finite number of errors produces the correct hypothesis ever after. Problems of this kind have an interesting complexity, and if the allowable semantic jumps must be permutations of some finite set there is a universal method that will solve any solvable discovery problem.

The discovery of *ceteris paribus* claims is not relativistic inquiry, but put with a learning theoretic framework there is a similar indexicality and a structural similarity to the relevant notions of convergence to the truth.

4.

I will give a more general form to the third ERS claim that putative *ceteris paribus* laws are not testable: there is no procedure for discovering the truth or falsity of *ceteris paribus* claims, even in the limit. My justification is that if there were some procedure that reliably, efficiently and feasibly led us to the truth or falsity of *ceteris paribus* claims but did not involve “testing”, the ERS claim would be, as Adolf Grunbaum likes to say, otiose, and if testing were not part of some procedure for getting to the truth, it would not serve the aims of inquiry.

I will interpret ERS as follows: A ***ceteris paribus* claim** has the form, *normally X*, or *ceteris paribus X*, where *X* contains no normalcy or *ceteris paribus* claim. An **instance** of *normally X* is a universal claim of the form *for all x, if A(x, z) then X(x)*, where *A* contains no normalcy claim but may have any quantifier structure. I will abbreviate this formulation as *if A then X*. *Ceteris paribus* generalizations, or claims that *normally, X*, are universal conditionals in which the antecedent is indexical and typically unexpressed; put another way, the apparent qualifier “normally” acts logically as a propositional function variable whose value is somehow determined in each case. *Ceteris paribus* generalizations are true for a case, a circumstance, a situation, if the appropriate value for the indexical antecedent for that case results in a true conditional, that is, if there is some appropriate instance of *normally X* that is true for the case. *Ceteris paribus*

generalizations are true if they are true for every case. The issue is what determines which instances are “appropriate.” I will pursue the idea that the learner – the inquirer – does the determining, although in all of what follows that determination could be made by any arbitrary community, e.g., the scientific community to which the learner belongs.

The speaker who endorses a *ceteris paribus* generalization endorses at least two things: there exist conditions A such that *if A then X*, and if the conditions he presently estimates to be such an A hold (his present normalcy conditions), then X holds of the case. The scientific or other interest of the *ceteris paribus* claim will of course turn on how frequently the normalcy conditions obtain, on how often they change, on how many logically independent propositions X share the same normalcy condition, and on other factors that are part of the routine of scientific assessment and debate.

The semantic oddity is that the *ceteris paribus* claim is indexical over the antecedent, that is, over the normalcy condition, but not over the consequent of the *ceteris paribus* claim, and, further, that the speaker, or inquirer, or learner, not the data, determine the normalcy condition for each case-but not its truth or falsity. In the course of inquiry into a *ceteris paribus* generalization the inquirer may alter his normalcy conditions for a proposition X – his current estimate of the antecedent to form the appropriate instance – while keeping the consequent, X, fixed, all the while endorsing the same tacit indexical conditional: *Normally, X*. (He may, of course, also abandon his endorsement of the *ceteris paribus* claim altogether.) In conjecturing *Normally, X*, the learner expresses a prediction about X and the conditions for normalcy that will obtain in the future. Learning “strict” laws is just the limiting case in which the normalcy condition is vacuous.

Less psychologically, and a bit more formally, a *ceteris paribus* learner for hypothesis X in the context of a particular discovery problem is a function that specifies, for each initial data segment, a consistent proposition A that (for simplicity) I will assume is bivalent for each initial data segment, and conjectures *if A then X* or alternatively, *if A then not X* for the next datum. (Recall that each datum may be a complex situation.) In the hands of the learner, the proposition *normally X* generates from the data a sequence of more specific conjectures: *if A1 then X, if A2 then X, . . . , if An then ~X, . . .* where some or all of the A_i and A_j may be the same proposition. The learner verifies that *normally X* for a data sequence if only a finite number of these conjectures is in error or is of the form *if An then ~X*, and falsifies that *normally X* for a data sequence if only a finite number of these conjectures is in error or is of the form *if An then X*. Reliable

verification or falsification is respectively verification or falsification over all data streams in the discovery problem.

5.

The formal picture is coherent, and it provides both a semantics and a methodology for *ceteris paribus* claims. *Normally X* is indexical because the content of “normally” depends on the inquirer, but the ideal inquirer’s sequence of conjectures, together with the data, determine whether *normally X* is true in any case, whether *normally X* is true in general, and whether normally *X* is true always after some time. Even so, one may wonder if the original doubts, formulated in ERS claims, have been addressed.

A second party, trying to determine whether a learner *L* can learn a claim, *normally X*, seems to be faced with a problem if *L* cannot specify, a priori, *L*’s function from finite data sequences to normalcy conditions for *X*. That may be an intuition behind the ERS doubts about *ceteris paribus* laws. They put the issues as semantic and methodological, but perhaps they are really about others and not about ourselves: how are we to know what someone means – what empirical conditions would make what is said true or make it false – when someone says “normally” *X*? But it would be a confusion to reject *ceteris paribus* laws on these grounds. Learning the learner’s function from data sequences to normalcy conditions for *X* is much harder, unnecessarily harder, than learning whether the learner converges to the correct truth value (or to the eventual truth value) about “normally *X*.” We can learn the latter in the limit (assuming, of course, that in each case the truth or falsity of the current normalcy condition for *X* and the truth or falsity of *X* can be determined). We simply conjecture “yes” when the current normalcy criterion yields a correct instance of “normally, *X*” and “no” otherwise.

The fact remains that some learner might, by having a policy of formulating stronger and stronger normalcy conditions, trivialize the content of the claim that *normally X*. Nothing in either the semantics or epistemology of *ceteris paribus* claims requires that a learner have such a policy, such a normalcy function, and when we reasonably suspect someone does we reasonably doubt his scientific seriousness. Epistemologically, the learning theoretic analysis of normalcy resembles in this respect the learning theoretic analysis of truth that is relative to the learner’s conjectures about the truth. There are consistent relativistic systems in which the learner’s conjecture that *X* makes *X* true, but these are not all, or even the most interesting, relativistic systems. Alternatively, of course, a learner might

be a fake who only pretends to a *ceteris paribus* hypothesis by observing whether *X* obtains in a case and also observing other features of the case and, if *X* obtains, announcing that conditions are normal, and, if *X* does not obtain, announcing a normality condition that he knows is not satisfied in the case. We have strategies for detecting such fakery, and they are imbedded in normal scientific practice.

6.

This preceding is my answer to the ERS objections to *ceteris paribus* laws, but it may be quite wrong. It may be that the real process of inquiry into many *ceteris paribus* laws about *X* aims to discover a fixed (but possibly disjunctive) set of features of the normal world sufficient for *X* and preferably almost necessary for *X*, or perhaps for many hypotheses simultaneously. If so, such *ceteris paribus* generalizations are simply place holders in the search for the kind of laws that ERS call strict. *Ceteris paribus*, *X*, is what we say when we believe that there are strict conditions *A* such that for all cases, if *A* then *X*, and we are able to specify *A* incompletely and believe we can identify normal conditions when they are present, but are disposed to investigate the features of normalcy further if counterexamples to *X* occur, and aim eventually to eliminate all normalcy qualifiers. I do not think this is so, but it may be.

A more likely alternative is that our actual strategy is sometimes to use counterexamples to generate new hypothesis, new *Xs*, that contain within them the conditions that apply when counterexamples occur. Perhaps *ceteris paribus* hypotheses are themselves part of a discovery strategy for new *ceteris paribus* hypotheses, and that is their special virtue. The genes of mouse liver cells express in a particular way on the surface of the Earth. Call the details *X*. Hammond discovered that in the microgravity aboard the Space Shuttle in Earth orbit, mouse kidney cells express their genes quite differently, call the details *Y*. Perhaps what we should now say is this: normally on Earth mouse kidney cells express as *X*, and normally in microgravity mouse kidney cells express as *Y*.

7.

So, I have more confidence that ERSs arguments are wrong than that my reconstruction of the semantics (ok, pragmatics) and methodology of *ceteris paribus* generalizations is right. In any case, the intuitions that I have guessed lie behind two of their three claims lead to an interesting

analysis with interesting possibilities, if not to their conclusions, and I am grateful for their provocation. There is a later, positive proposal in their paper for which I am not grateful, and I cannot let it pass without comment.

Explicitly discussing causal claims, such as “smoking causes cancer” ERS say that this: they have no content, they say nothing, they merely “signal” or “express” the empirical data that prompts their assertion, and perhaps also signal a “research program” to “explain” their consequents in terms of their antecedents. Taking the proposal seriously, the introduction and discussion sections of most articles in *Science*, and all articles in *Cell*, should be excised, or at least retitled “Signaling.” One can imagine the conversations:

Russian scientist who knows only of Russian data on smoking and cancer:

“Smoking causes cancer.”

An oncologist, who does not know of Russian data and has a different research program:

“Smoking causes cancer.”

Russian scientist:

“We agree, then.”

Earman, Roberts and Smith:

“No you don’t.”

Meanwhile we can reflect on the predicament of the oncologist who, saying “smoking causes lung cancer”, signals her program to “explain” lung cancer from smoking while the content of her explanations of lung cancer from smoking, beyond references to data, are nothing but more signals that she intends to carry out research to “explain” lung cancer from smoking. And on and on, forever.

In a time in which many first class minds are simultaneously struggling towards an understanding of causation and reliable methods for determining causal relations in complex circumstances, we should not be grateful for this.

REFERENCES

Kelly, K. T.: 1995, *The Logic of Reliable Inquiry*, Oxford University Press, Oxford.

- Kelly, K.: 1998, 'Iterated Belief Revision, Reliability, and Inductive Amnesia', *Erkenntnis* **50**, 11–58.
- Kelly, K. T. and C. Glymour: 1992, 'Inductive Inference from Theory Laden Data', *Journal of Philosophical Logic* **21**, 391–444.
- Kelly, K., C. Juhl, and C. Glymour: 1994, 'Reliability, Realism, and Relativism', in P. Clark (ed.), *Reading Putnam*, Blackwell, London, pp. 98–161.
- Kitcher, P.: 1990, 'The Division of Cognitive Labor', *Journal of Philosophy* **87**, 5–22.
- Kitcher, P.: 1995, *Advancement of Science: Science without Legend, Objectivity without Illusions*, Oxford University Press, Oxford.

Clark Glymour
Department of Philosophy
135 Baker Hall
Carnegie Mellon University
Pittsburgh, PA 15213
U.S.A.
Email: cg09@andrew.cmu.edu
Institute for Human and Machine Cognition,
University of West Florida

