
Economic Models

Author(s): Allan Gibbard and Hal R. Varian

Source: *The Journal of Philosophy*, Nov., 1978, Vol. 75, No. 11 (Nov., 1978), pp. 664-677

Published by: Journal of Philosophy, Inc.

Stable URL: <https://www.jstor.org/stable/2025484>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*

JSTOR

guage, since there are serious difficulties in describing the properties languages are *known* to have, and since there are still greater difficulties in explaining why they should have *these* properties and not others, it seems to me that whether propositions exist is itself an unclear question.

JAMES HIGGINBOTHAM

Columbia University

ECONOMIC MODELS *

IN the past few decades, model building in economics has yielded many powerful, clear results with rigorous demonstrations. In some cases, the new results explicate old conclusions which were vaguely formulated and supported. According to Adam Smith, every individual, without intending it, "labours to render the annual revenue of the society as great as he can."¹ Beginning in the 1930s, models were constructed in which this claim of Smith's was precisely formulated, and conditions were found under which the claim would be true.² These conditions include perfect rationality and perfect information on the part of economic agents, divisibility of all factors of production, privacy of all goods, and nonincreasing returns to scale for each factor. More recent models have addressed limitations on information, relaxations of the standard rationality conditions for economic agents, public goods and externalities, and other deviations from the model of perfect competition.

These triumphs of economic model building can evoke suspicion. Do they tell us anything about the world of genuine people, work, production, and commerce? The assumptions of the models often do not seem even remotely accurate as descriptions of an actual economy, and the happy consequences of some models of

* To be presented in an APA symposium on the Philosophy of Economics, December 29, 1978. Russell Dancy and Alexander Rosenberg will comment; see this JOURNAL, this issue, 677-679 and 679-683, respectively.

¹ *The Wealth of Nations* (1776), E. Canaan, ed. (New York: The Modern Library, 1937), p. 423.

² For a brief history of this, see Kenneth J. Arrow and F. H. Hahn, *General Competitive Analysis* (New York: Holden Day, 1971), ch. 1. Expositions of the subject are to be found in Tjalling C. Koopmans, *Three Essays on the State of Economic Science* (New York: McGraw-Hill, 1957), Essay I, and Hal R. Varian, *Microeconomic Analysis* (New York: Norton, 1978), ch. 5.

perfect competition do not match our experience. Perhaps economic models, despite their apparent power, turn out to be mere exercises in mathematics or apologies for *laissez faire* capitalism.³

Can models with unrealistic assumptions, then, be of any use in understanding the world? We think that they can be, and in this paper, we shall discuss how. The issue, of course, needs a more precise formulation. Assumptions cannot meaningfully be called “unrealistic” without more said, for the same model can be applied to many different situations in the world, and its assumptions may be realistic for some of these situations and not for others. The question we should be asking is: In what ways can a model help in understanding a situation in the world when its assumptions, as applied to that situation, are false?

Our emphasis here will be on the use of models by that group of economists known as “economic theorists.” Large-scale econometric models programmed on computers are the major tools for forecasting the performance of an economy, but about them we shall have little to say. Within the class of theoretical models, we can distinguish between “descriptive” and “ideal” models. Descriptive models attempt to describe, *in some sense*, economic reality. Ideal models, on the other hand, are concerned with the description of some ideal case which is interesting either in its own right or by comparison to reality. Our emphasis here will be primarily on descriptive models.

Finally, within this subclass of descriptive models, we shall distinguish between models that are “approximations” and models that are “caricatures.” The former are models that aim to describe reality, albeit in an approximate way. Caricatures, on the other hand, seek to “give an impression” of some aspect of economic reality not by describing it directly, but rather by emphasizing—even to the point of distorting—certain selected aspects of the economic situation.

We might liken econometric models, approximations, and caricatures, to photographs, realistic drawings, and caricatures in the literal sense. For some applications, a photograph may be the best means of depicting an object, but in some cases a drawing or even a caricature may allow greater understanding. Furthermore, a commercial artist who is attempting to provide an accurate impression

³ Two recent critiques of present neoclassical economic practice are Benjamin Ward, *What's Wrong with Economics* (New York: Basic Books, 1972), and Martin Hollis and Edward Nell, *Rational Economic Man* (London: Cambridge, 1975).

of some object may utilize all three media, either sequentially or simultaneously. It is the same with the three techniques of economic modeling.

One way in which our terminology is misleading is that it suggests clear, separate categories of models, whereas the distinction is one of degree. In many cases, an economic phenomenon will initially be represented by a caricature, and the representation will then gradually evolve into an econometrically estimable model.

In our discussion of the use of models by economic theorists, we shall ignore some important issues. How significant is the understanding yielded by models likely to be: is the pursuit of model building the most promising road to understanding the economic world? In what ways is economics like the natural sciences in its use of models, and in what ways is it different? These are questions we shall skirt; our concern will simply be with how, if at all, models with assumptions false of an economic situation can help us to understand that situation.

I. A CHARACTERIZATION OF ECONOMIC MODELS

As we are using the term, a *model* is involved whenever there is economic reasoning from exactly specified premises. In not all such cases do economists use the term 'model'; what are standardly called "the theory of the firm" and "the theory of the consumer" involve models in the sense in which we are using the term.⁴

A model, we shall say, is a *story* with a specified *structure*: to explain this catch phrase is to explain what a model is. The *structure* is given by the logical and mathematical form of a set of postulates, the *assumptions* of the model. The structure forms an uninterpreted system, in much the way the postulates of a pure geometry are now commonly regarded as doing. The theorems that follow from the postulates tell us things about the structure that may not be apparent from an examination of the postulates alone.

Although the term 'model' is often applied to a structure alone, we shall use it in another sense. In economists' use of models, there is always an element of interpretation: the model always tells a story. If we think of the structure as containing uninterpreted predicates, quantifiers, and the like, we can think of the story as telling what kind of extension each predicate has and what kind of domain each quantifier has: a model will talk of firms, consumers, preferences, prices, information, and the like. The story may be vague: a model involving preferences and information, for

⁴ For many examples of economic models, see Varian, *op. cit.*

instance, will ordinarily offer no explication of what preferences and information are, beyond what it says about their structure. The structure itself must be specified with the precision needed for mathematical reasoning.

Sometimes it will be found that two models, with two different stories, have the same structure, or that one model has the structure of the other and some additional structure as well. A dynamic model, for instance, may turn out to have the structure of a static model, although the story of the dynamic model involves time and that of the static model does not. When a structure is shared, theorems about one model—which, after all, characterize its structure—carry over when the story of the other model is told. This reinterpretation of structures is ubiquitous in economic theorizing, and is often a source of great insight.⁵

Now although a model, as we use the term, is a story, it is not a story about any particular situation in the world. The theory of the firm, for instance, does not tell us which firms it is describing. The assumptions and derived statements of a model, then, are not themselves propositions that can be true or false, roughly true or wildly off the mark.

A model, though, may be applied to the world or, as we shall say, applied to a *situation*. It is then said what firms, what prices, and the like are being modeled; the result we shall call an *applied model*. An applied model is stipulated by starting with a model and then giving its predicates particular extensions, its quantifiers particular domains, and the like—by providing an “interpretation” in the logicians’ sense. At least that is what would happen if applied models were built with logicians’ tools; in practice, the application of a model is likely to be informal, and it will often be vague or even implicit. However an applied model is presented, the difference between a model and an applied model is this: whereas a model speaks of entities of certain general kinds—prices, consumers, information, and the like—without saying which particular entities in the world they are, an applied model specifies the particular classes of entities it treats. The theory of the firm, for instance, is a model; when it is interpreted as talking about General Motors, the cars General Motors produces, their prices, and the like, it is an applied model.

The distinction between a model and an applied model will be

⁵ Where we speak of different models sharing an interpretation, Koopmans (60–62) speaks of different interpretations of the same model. He gives a number of examples.

central to our discussion. For whereas the statements of a model are not propositions that are true, nearly true, or grossly false, the statements of an applied model are. When a model is applied to a situation, we can ask how close to the truth its statements are.

In our present vocabulary, the topic of this paper is how a model can be helpful in understanding a situation if, when the model is applied to the situation, the assumptions of the resulting applied model are false. One way to approach this question is to consider the relation between the assumptions of an applied model and hypotheses about a situation. By a *hypothesis*, we mean an empirical proposition that an investigator thinks sufficiently plausible to make it worth while to evaluate the evidence for its truth and falsity. Are the assumptions of an applied model, we may ask, hypotheses about the situation to which it is applied? Is there some other uniform relation between the statements of an applied model and the hypotheses of the investigator who uses the model to understand a situation?

II. MODELS AS APPROXIMATIONS

All economic models have this, at least, in common: a model poses a question of the form, "What would happen if such and such were the case?" in such a way that it can be answered deductively. Perhaps economic models have more in common than this, but this characterization provides a starting point for an investigation. We can now ask why questions of this form might matter. What do the answers to such questions tell us about economic situations? In asking this, of course, we need not suppose that answers to all such questions are equally helpful. It may be that answers to some such questions will help us in understanding an economy, and answers to others will not. If so, we need to ask what sorts of models are useful in economics and what sorts are not, and why.

Knowing what would happen if such and such were the case is, of course, useful if we think that perhaps it *is* the case. An investigator who applies a model to a situation might hypothesize that the assumptions of the applied model are true of the situation. That suggests a theory of economic models which we shall dismiss as the *naive view*: that when an investigator legitimately applies a model to a situation, he investigates the hypothesis that the assumptions of his applied model are true of the situation.

Why do we dismiss this view as naive? In the first place, it travesties economic practice. When an economic theorist applies a model to a situation, it is almost always preposterous to suppose that the assumptions of the applied model are exactly true of the

situation. Indeed the only statements of most applied models in economics that are true exactly are truths with no empirical content, such as definitions and mathematical truths. We should look for an account of economic models that makes sense of what economists do with them. In the second place, when practicing economists have discussed the use of models, they have not adopted the view we label "naive." They have not supposed that the assumptions of an applied model themselves constitute a hypothesis about the situation to which the model is applied. Rather, they talk of the assumptions of a successfully applied model as approximations. Put in our language, the prevailing view is that, when an investigator applies a model to a situation, he hypothesizes that the assumptions of the applied model are close enough to the truth for his purposes.⁶

This answer must be right in its essentials, at least for many uses of models. More, though, needs to be said. The rough truth of a model is supposed to explain something about a situation; how can it do so?

The first step in the explanation is to transform what is to be explained. When a model is used as an explanatory approximation, it is claimed that certain propositions of the applied model are roughly true; what is now to be explained is the fact of their approximate truth. Call these propositions the *conclusions* of the applied model. There will be no uniform relation between this new *explicandum*—the approximate truth of the conclusions of the applied model—and what was to be explained before the model was applied. Perhaps what was originally to be explained was formulated independently of the model, and to explain the approximate truth of the conclusions of the model is at least partially to explain the original *explicandum*. Perhaps it is initially unclear what is to be explained, and a model provides a means of formulation. Perhaps a statement is derived from the assumptions of a model before the investigator even thinks about its applicability; afterwards, he may apply the model to a situation, hypothesize that the statement as applied is roughly true, and seek to explain its rough truth by using the model. In all these cases, what

⁶ See, for instance, Koopmans, *op. cit.*, pp. 142–144, and Abraham Wald, "On Some Systems of Equations of Mathematical Economics," *Econometrica*, xix (1951): 369. (This is a translation of his pioneering 1936 article). For a variant of this view which we discuss briefly later, see Milton Friedman, "The Methodology of Positive Economics," in *Essays in Positive Economics* (Chicago: University Press, 1953).

the application of the model is to explain is formulated by means of the model.

An explanation by the approximate truth of a model takes the following form. First, if the assumptions of the applied model were *true*, the conclusions would be—here the proof is mathematical. Second, the assumptions in fact are sufficiently close to the truth to make the conclusions approximately true. For this no argument within the model can be given; it is rather a hypothesis, for and against which evidence might be given. One kind of evidence is evidence for the rough truth of the conclusions of the applied model; another kind is evidence for the rough truth of its assumptions. A third kind of evidence might employ a new model. Suppose the assumptions of an applied model, we have reason to think, deviate from the truth in a systematic way. We might test that model by another applied model whose assumptions, we think, more closely approximate the truth. We can then see whether differences in the assumptions of the two models make a significant difference to their conclusions. If not, that is evidence for the hypothesis that the conclusions of the original applied model were close to the truth *because* its assumptions were sufficiently close to the truth. When we vary the assumptions of a model in this way to see how the conclusions change, we might say we are examining the *robustness* of the model. We shall have more to say about robustness later.

Another kind of evidence might appeal to the nature of the deviations from the model's assumptions. If deviations are random, or more precisely, are not systematic, there might be good reason to have some faith in the conclusions of the model even though the assumptions, strictly interpreted, are implausible. Perhaps a case in point is the economist's assumption of perfect optimizing behavior. Of course this assumption is, strictly speaking, false, but, so long as errors in optimization are not systematic, this hypothesis may be useful in describing the "central tendency" of economic behavior. Furthermore, in models where individual units' behavior is being aggregated, nonsystematic errors may be expected to "wash out" in the process of aggregation. We appeal here to no general principle that nonsystematic errors in the assumptions of an applied model leave its conclusions untouched,⁷ but in many cases that can be expected to happen.

⁷ It might happen, for instance, that businesses were able to exploit non-systematic deviations from perfect optimizing behavior—to profit from random suckers—and that that exploitation has significant economic effects.

Milton Friedman gives an account of economic theories which, on a reading that fits most of what he says, differs from the account we have given. On this reading of Friedman, when a model is applied to a situation, all that is hypothesized is that the conclusions of the applied model are close enough to the truth for the purpose at hand.⁸ According to us, something further is hypothesized: that the conclusions are sufficiently close to the truth *because* the assumptions are sufficiently close to the truth. Friedman's reason for taking the position he does appears at least in part to be as follows. The assumptions of a model may approximate reality sufficiently for some purposes but not for others; the assumption of a vacuum, for instance, approximates reality sufficiently in an explanation of the rate of fall of a compact ball from a roof, but not in an explanation of the rate of fall of a feather. We have no standard, then, for when the assumptions of an applied model are sufficiently realistic, and so we should judge an applied model entirely by the accuracy of its conclusions.

A serious problem with Friedman's position is that economic models are often used to extrapolate to new situations. We are often interested in the effect on economic behavior of changing the economic environment. In order to have any faith at all in such extrapolations, we must believe that there is some sort of connection between the accuracy of the assumptions and the accuracy of the conclusions. Our view accepts part of Friedman's reasoning on the matter: there is, we agree, no standard independent of the accuracy of the conclusions of an applied model for when its assumptions are sufficiently realistic. There is, though, a *derived* standard for the accuracy of the assumptions. If accuracy or degree of approximation were numerically characterizable, our position could be put as follows. When a model is applied to a situation as an approximation, an aspiration level ϵ is set for the degree of approximation of the conclusions. What is hypothesized is this: there is a δ such that (i) the assumptions of the applied

⁸ *Op. cit.*, pp. 7–30. A theory, says Friedman, “is to be judged by its predictive power for the class of phenomena it is intended to ‘explain’” (8). The view that the conformity of a theory's “assumptions” to “reality” provides an additional test “is fundamentally wrong and productive of much mischief” (14). Some of the things Friedman says, though, fit the view we are advocating. “To put this point less paradoxically, the relevant question to ask about the ‘assumptions’ of a theory is not whether they are descriptively ‘realistic,’ for they never are, but whether they are sufficiently good approximations for the purpose in hand” (15). For good discussions of Friedman's views, see Alexander Rosenberg, *Microeconomic Laws* (Pittsburgh, Pa.: University Press, 1976), pp. 155–170, and Koopmans, 137–140.

model are true to degree of approximation δ , and (ii) in any possible situation to which the model could be applied, if the assumptions of that applied model were true to degree of approximation δ , its conclusions would be true to degree ϵ . Of course, when models are applied as approximations, few if any of the degrees of approximation involved are characterized numerically, but the pattern of explanation is, we think, the one we have given.

III. THE FIT OF MODELS TO THE WORLD

Economists apply models to situations in two quite different ways, which we shall call *econometrically* and *casually*. The purposes of the two sorts of applications are different, and the kinds of models that are applied in the two ways are different.

The goal of casual application is to explain aspects of the world that can be noticed or conjectured without explicit techniques of measurement. In some cases, an aspect of the world (such as price dispersal, housing segregation, and the like) is noticed, and certain aspects of the micro-situation are thought perhaps to explain it; a model is then constructed to provide the explanation. In other cases, an aspect of the micro-world is noticed, and a model is used to investigate the kinds of effects such a factor could be expected to have. If the model turns out to have striking features, a casual search for economic situations with those features may then be conducted. In either kind of case, no measurement that goes beyond casual observation is involved.

When economic models are used in this way to explain casually observable features of the world, it is important that one be able to grasp the explanation. Simplicity, then, will be a highly desirable feature of such models. Complications to get as close as possible a fit to reality will be undesirable if they make the model less possible to grasp. Such complications may, moreover, be unnecessary, since the aspects of the world the model is used to explain are not precisely measured.

About econometric applications we shall say little; there is a well-developed methodology of the subject.⁹ Such applications require explicit techniques of measurement for some of the quantities in the model, though ordinarily not for all. Since the goal is to achieve a close fit with measured reality, complexities that help to achieve this close fit may be tolerated. We should expect, then, that models intended for econometric application will be more complex than those intended for casual application. Highly com-

⁹ See Varian, *op. cit.*, ch. 4, and the references cited there.

plex econometric models often are intended more for prediction than for explanation, and, when they are used to explain, it may be through quasi-experimentation by varying parameters. On the other hand, econometric models have for the most part evolved from simpler approximation models. Indeed, it is generally consideration of the structure of such simpler models that determines what measurements and statistical techniques are relevant for formulating and estimating econometric models.

IV. MODELS AS CARICATURES

What can clearly be said of any model, we noted at the outset, is that it poses a question "What would happen if such and such were the case?" so that it can be answered deductively. We then asked how questions of that form might be valuable for understanding economic situations. One way, which we have been discussing, is through approximation: one can sometimes usefully hypothesize that the "such and such" approximates reality to a degree sufficient for a purpose.

Now although that may be the most important use of models in economics, it is by no means their only important use. Often the assumptions of a model are chosen not to approximate reality, but to exaggerate or isolate some feature of reality. An applied model that ascribed to that feature its approximate place in reality might bury its effects, and, for that reason, a model that is a better approximation to reality may make for a worse explanation of the role of some particular feature of reality.

If the purpose of economic models were simply to approximate reality in a tractable way, then, as techniques for dealing with models are refined and as more complex models become tractable, we should expect a tendency toward a better fit with complex reality through more and more complex models. That sometimes happens, but a tendency to better approximations through more complex models is by no means the rule. Often, a feature of the world that might have been added to a model as a complication is instead treated as a central aspect of a new, simple model.

When that happens, the representation of the feature is not so much an approximate description of the feature and its place in the world as it is a caricature. By that we mean not only that the approximation is rough and simple, but that the degree of approximation is not an important consideration in the design of the model. What typically happens is this. First, some aspect of economic life is noted. (For instance, learning which store charges least for an item takes effort, and some consumers spend the effort

whereas others do not. Stores use strategies to make money from customers of both kinds. Sales and low prices attract informed customers; high prices make high unit profits from uninformed customers.) The theorist sets out to construct an explanatory model. (How much of stores' pricing policies, he may ask, can be explained on the basis of the factors noted?) The reality he confronts is complex. (No simple, accurate account could be given of the forms consumers' information takes, or of their motives in seeking information, choosing stores, and making purchases.) The theorist's approach to his problem is not to try for the closest tractable approximation of this complex reality, but to tell a simple story that captures some of its features. (Perhaps in the story, the consumer's choice is simply this: he can buy the information of which store charges the lowest price, or he can choose a store at random. Firms adopt pricing strategies that maximize their profits.) The model poses the question, what would happen if this story were true?¹⁰

When can such a caricature be helpful in understanding a situation? One way is by yielding conclusions that are *robust*, in the sense that they do not depend on the details of the assumptions.¹¹ When a theorist applies a model that caricatures a situation, one hypothesis he may entertain is this: the conclusions of

¹⁰ The caricature sketched in parentheses is from S. Salop and J. Stiglitz, "Bargains and Ripoffs", *Review of Economic Studies*, XLIV (1976): 493-510. Here are two other examples of caricatures in economics. First, Paul A. Samuelson, "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money," *The Collected Papers of Paul A. Samuelson*, vol. I, (Cambridge, Mass.: MIT Press, 1966). Samuelson is concerned with modeling the rationale behind intergenerational transfers of income, such as social security programs. He assumes a population of exactly similar people, each of whom lives only two periods, working one period and retiring the next. This is obviously a gross distortion of reality, but, by means of it, Samuelson is able to shed some light on the virtues and limitations of such programs. Second, M. Spence, *Market Signalling* (Cambridge, Mass.: Harvard, 1974). An employer may require a college degree for a job not only because he thinks that students acquire skills in college that are useful in that job, but because he thinks that the sort of person who is likely to attain a degree is the sort of person who will do the job well. That may be part of the reason that people with college degrees command higher salaries, as a rule, than people without degrees, and to the extent that it is, people may have an incentive to spend on amount of time and effort on college that is wasteful. To investigate this situation, Spence starts out with a model in which education has no influence on worker productivity at all. He assumes this not because he supposes it to be true, approximately true, or even the best approximation to the truth of its degree of simplicity, but rather because he can then more clearly isolate the effects he wishes to model.

¹¹ Arrow and Hahn use the word in this sense, *op. cit.*, p. vii. Statisticians speak of the "robustness of an estimator" in an analogous sense.

the applied model roughly depict some feature of the situation, and that is because (1) the assumptions of the model caricature features of the situation, and (2) the conclusions are robust under changes in the caricature. A principal way of testing this hypothesis may be to try out models with disparate caricatures of the same complex aspect of reality.

A striking feature of microeconomic explanation is this. Individuals, acting from familiar motives, can interact to produce large-scale economic effects that could not be expected at all on the basis of a naive inspection of the individuals' dispositions. Of an alleged happy instance of this, Adam Smith said, the individual "is led by an invisible hand to promote an end which was no part of his intention" (423). Thomas Shelling speaks of "the ecology of micromotives."¹² To the extent that macrophenomena can be explained through an ecology of micromotives, it must be possible to find in the complex interplay of micromotives a pattern or tendency on which the macrophenomena depend. Now that, to be sure, can be done in a number of ways: sometimes statistical characterization of the micromotives will identify the required tendency, and sometimes it will be possible to give an approximate description of the micromotives, with the claim that the truth deviates from the approximation in no way that systematically affects the conclusions. Often, though, macrophenomena can be explained on the basis of the interplay of micromotives only because they are robust under variations in the micromotives. Probably in much of economics, the macrophenomena that can be explained are those which are robust under variations in individual dispositions. One way to test for such robustness is to look for conclusions of models that are robust under different caricatures of a feature of reality.¹³ In that search, the degree of approximation to the truth achieved by any particular caricature is beside the point.

Is a caricature no more than an approximation that is especially rough and simple? What we have said about approximations seems to apply to caricatures. When an applied model is used as an ap-

¹² "On the Ecology of Micromotives", *The Public Interest*, xxv (1971): 59-98.

¹³ For example, the standard economic model of firm behavior is a model of perfect profit maximization. This obviously exaggerates the acumen of real-world managers. However, there is a class of "evolutionary" models in which firms behave as imperfect optimizers, but considerations of survival of the fittest turn out to imply that the long-run behavior of an industry composed of such firms is similar to that of an industry composed of perfect optimizers. For examples of models of this sort, see Richard H. Day and Theodore Groves, eds., *Symposium on Adaptive Economics*, 1 (New York: Academic Press, 1975).

proximation, we have said, the investigator sets an aspiration level for the accuracy of the conclusions. He hypothesizes that the assumptions are near enough to the truth for the conclusions to achieve that level of accuracy—and indeed for them to achieve it *because* the assumptions are sufficiently close to the truth. Perhaps when an applied model is used as a caricature, that just means that the aspiration level is low, so that various extremely simple models achieve it, and indeed do so because only the roughest fit of assumptions to situation is needed to achieve it.

A caricature differs from an approximation, though, not only in its simplicity and inaccuracy, but in its deliberate distortion of reality. When a model is applied as an approximation, the goal is to distort as little as is compatible with a given degree of simplicity and tractability. A caricature involves deliberate distortion for other reasons—to isolate the effects of one of the factors involved in the situation, or to test for robustness under changes of caricature. Although, then, when a model is applied as a caricature, it may indeed be hypothesized that the model is an approximation of the roughest kind, the model will be chosen not for the sake of good approximation, but to distort reality in a way that illuminates certain aspects of that reality. That, of course, is what a pictorial caricature does.

VI. CONCLUDING REMARKS

Much of economic theorizing consists not of an overt search for economic laws, not of forming explicit hypotheses about situations and testing them, but of investigating economic models. A model always poses a question, What would happen if such and such were the case? What, we have asked, can an answer to this question tell us of economic life? Often when a model is presented, only the briefest suggestive remarks are made about its bearing on the world, and yet it seems clear that, when an economist investigates a model, it is often because he thinks the model will help to explain something about the world. We have discussed two patterns of explanation—two patterns of what an investigator's hypothesis about a situation may be when he applies a model to it. The hypothesis may be that the conclusions of an applied model are approximately true, and that that is because its assumptions are sufficiently close to the truth. In some such cases, the hypothesis is tested casually; in others, econometrically; quite different kinds of models lend themselves to the two kinds of testing. The hypothesis may, on the other hand, be that a conclusion of the applied model depicts a tendency of the situation, and that this is because

the assumptions caricature features of the situation and the conclusion is robust under changes of caricature. The distinction between the two kinds of hypotheses is not sharp. A caricature may be, among other things, an approximation of a particularly rough and simple kind. The difference between applying a model as an approximation and applying it as a caricature lies in the intentions of the investigator: a caricature involves deliberate distortion to illuminate an aspect of economic life. If the uses of deliberate distortion are ignored, and the job of applied models is taken to be no more than accurate approximation under constraints of simplicity and tractability, many of the caricatures economic theorists construct will seem unsuited for their job.

ALLAN GIBBARD

HAL R. VARIAN

University of Michigan

MODEL BEHAVIOR *

IF Gibbard and Varian are right, who is wrong?

In an early footnote they refer to Hollis and Nell (*Rational Economic Man*; for publishing data on this and other references below, see Gibbard and Varian's notes). But these latter authors end their book with an elaborate presentation of models along Ricardian-Marxist-Sraffish lines. They do not need to be told that economic models are at best approximations, or highly selective. But that is what Gibbard and Varian tell them, in the expectation that it will have some bearing on the feeling that Hollis and Nell share with many of the rest of us, that current micro-economic models may be (in Gibbard and Varian's words) "mere exercises in mathematics or apologies for *laissez-faire* capitalism."

But Gibbard and Varian's message has no bearing on that feeling, for those who share it are rejecting neo-classical theory not because it employs models, but because it employs poor models, and pronouncing them poor not because they are only approximations, or highly selective, but because they are not even approximations, and not selective but fictional.

* Abstract of a paper to be presented in an APA symposium on the Philosophy of Economics, December 29, 1978, commenting on Allan Gibbard and Hal R. Varian, "Economic Models"; see this JOURNAL, this issue, 664-677.

0022-362X/78/7511/0677\$00.50

© 1978 The Journal of Philosophy, Inc.