Economic Models as Analogies *

Itzhak Gilboa[†], Andrew Postlewaite[‡],
Larry Samuelson,[§] and David Schmeidler[¶]

This version, January 27, 2013

Abstract

People often wonder why economists analyze models whose assumptions are known to be false, while economists feel that they learn a great deal from such exercises. We suggest that part of the knowledge generated by academic economists is case-based rather than rule-based. That is, instead of offering general rules or theories that should be contrasted with data, economists often analyze models that are "theoretical cases", which help understand economic problems by drawing analogies between the model and the problem. According to this view, economic models, empirical data, experimental results and other sources of knowledge are all on equal footing, that is, they all provide cases to which a given problem can be compared. We offer complexity arguments that explain why case-based reasoning may sometimes be the method of choice and why economists prefer simple cases.

*We are thankful to many colleagues and friends with whom we have discussed the issues addressed here over the years. We thank Hervé Crès, Robin Cubitt, Eddie Dekel, Brian Hill, Doron Ravid, Jack Vromen, Bernard Walliser, three referees and the editor for comments on earlier drafts of this paper. We also thank Daria Engel and Kartin Kish for their research assistance. We gratefully acknowledge ISF Grant 396/10 (Gilboa and Schmeidler), ERC Grant 269754 (Gilboa), and NSF Grants SES-0961540 (Postlewaite) and SES-0850263 (Samuelson).

†HEC, Paris, and Tel-Aviv University, tzachigilboa@gmail.com.

‡University of Pennsylvania. apostlew@econ.sas.upenn.edu.

§Department of Economics, Yale University, Larry.Samuelson@yale.edu.

¶The InterDisciplinary Center in Herzliya, and TAU, davidschmeidler@gmail.com.

Economic Models as Analogies

Contents

Economic Models as Analogies

1 Introduction

Many economists think of their discipline as a successful social science. Economics relies on rigorous and non-trivial mathematical and statistical analyses. The paradigm of microeconomics is viewed as a unified approach that can deal with all problems of social interaction, and it is indeed adopted by other disciplines. Economics is popular with students, and economics professors are in high demand within the academic world and outside it.[1] However, the basic assumptions of economic theories have been harshly criticized by psychologists, presumably showing in laboratory experiments that these assumptions frequently do not hold.

If the assumptions of economics are all wrong, why do economists keep using them? Why do they develop sophisticated mathematical models based on such flimsy foundations? This question is our starting point. We elaborate on it and discuss two additional puzzles in Section 2. We then offer a possible resolution to these puzzles in Section 3, namely that some of the reasoning in economics is case-based rather than rule-based, and that economists view themselves as generating the "cases" to which real problems might be analogous. True to the method of our discipline, we construct a model (in Section 5) that illustrates the advantages of this mode of research. In Section 6 we discuss the virtues of a general paradigm, or, to be precise, of the standard language that such a paradigm employs. Section 7 concludes.

[1]Throughout this paper we make various claims about the sociology of economics. We draw on our personal impressions of the field over several decades, based on many discussions with colleagues, editorial work, and the like, but our casual observations are not backed by any scientific data.

2 Puzzles in the Sociology of Economics

In this section we describe three puzzles that, we will later argue, may be explained by understanding the way economists think about models. The questions we raise and answers we provide are descriptive rather than normative. We do not make here any claims about the optimality of the method of research that economics has adopted.

2.1 Assumptions are False

That the assumptions of economics are false is one of the most poorly kept secrets in science. Already in the 1950s, Milton Friedman felt that the issue was important enough to deserve a serious treatment. Friedman (1953) made the claim that economists should not worry if their assumptions (on individual behavior) are wrong, as long as their conclusions (regarding market phenomena) are relatively accurate. Friedman's defense came under various attacks, which we will not review. We only mention that microeconomics has changed its focus since Friedman proposed his defense, with more instances of individual behavior now considered part of the domain of economics than in the past, rendering the defense more problematic.

In the early 1970s, Daniel Kahneman and Amos Tversky launched a decades-long project that is sometimes summarized as "proving that people are irrational". Amos Tversky used to say, "Give me an axiom [on individual behavior] and I'll design the experiment that refutes it". Indeed, the psychological literature today is replete with examples of such experiments. After several decades in which economics essentially ignored the Kahneman-Tversky project, change began to appear in the mid-1990s. Behavioral economics has since been developed, making economic models more realistic by modifying them to be consistent with psychological findings. Many economists remain skeptical about the field, despite such recognition as the 2002 Nobel Prize awarded to Daniel Kahneman, but not because they believe that the classical assumptions are literally true.

Why does economic theory engage in relatively heavy technical analysis, when its basic premises are so inaccurate? Given the various violations of fundamental economic assumptions in psychological experiments, what is the point in deriving elaborate and carefully proved deductions from these assumptions? Why do economists believe that they learn something useful from analyzing models that are based on wrong assumptions?

2.2 Mathematization

A scientific field can sometimes be reduced, at least in principle, to another. Chemistry is, in principle, reducible to physics, biology to chemistry, and psychology to biology. By the same token, the social sciences, namely, economics, sociology, and political science, are in principle reducible to psychology. Of course, these reductions are highly theoretical and no one would seriously suggest that the behavior of countries should be analyzed by studying the motion of elementary particles. Yet, it is often useful to think in terms of the reliance of one scientific domain on another.

One typically finds a heavier reliance on mathematical analysis as one moves down the reduction chain. Physics is inarguably the most mathematized field, chemistry is less mathematical, and so forth. However, economics seems to be an exception to this rule. Economics engages in mathematical analysis that appears in general to be more sophisticated than that employed by psychology or even biology.

There is no a priori necessity that more basic fields will be more mathematized than the fields that rely on them. However, the apparent exception provided by economics and psychology may serve as a hint that economists think of their mathematical models differently than do other scientists.

2.3 The Scope of Models

Daniel Kahneman once noted[2] that psychologists and economists treat models very differently: psychologists are careful to define the scope of applicability of their models very precisely. Trying to avoid refutations of their theories, or failure to reproduce their findings, they seek a narrow definition of the applicability of the model or the theory in question. By contrast, he argued, economists tend to find their models useful in a wide variety of examples, viewing the latter as special cases of their model. Why do the two disciplines use models so differently?

2.4 Related Literature

A widely accepted observation is that the use of models in modern economic theory is sometimes quite different from its use in other sciences, as well as in some sub-fields of economics itself. Indeed, our casual sampling of colleagues and coauthors suggests that most economic theorists found it necessary to discuss the methodology of economic modeling in classes and in interaction with scholars from other disciplines.

Many economists and philosophers have also written on this topic. While we do not provide here an exhaustive survey of the philosophical and methodological literature on the topic, we mention several contributions.

Gibbard and Varian (1978) likened economic models to paintings, drawing, and caricatures. They argued that there are economic models that are supposed to mimic reality, as do paintings; others are supposed to simplify reality as do drawings; and yet others are meant to be exaggerated and distorted depictions of reality, as are caricatures. Hausman (1992) pointed out that economic theory models differ from econometric models, and that the former can be viewed as explorations. Maki (1994) highlighted the role of modelling as isolation, and more recently (Maki (2005)) argued that mod-

[2] At a talk at the Cowles Foundation, Yale University, September 2001.

els can be viewed as (thought) experiments (as well as that experiments can be viewed as models). Cartwright (1998) views models as attempting to establish capacities, which specify that under certain conditions, certain conclusions can be inferred.

Sugden (2000, 2009, 2011) discussed models in economic theory (as well as in mathematical biology) as "credible worlds" that are used to reason about reality. According to his view, a model does not attempt to describe the real world, but rather, a parallel, hypothetical world. If this model is "credible", it can be used to reason about the real world by means of inductive inference. He refers to analogies and similarities in his account of the use of credible worlds, but also to induction and abduction. Thus, his terms "inductive inference" seems to refer both to case-to-rule and to case-to-case induction. Sugden argued that economic models often lack a direct motivation in terms of an unexplained phenomenon, and, importantly, also concrete guidance about the applicability of these models. Thus, he holds, inductive inference from economic models depends on subjective judgments of similarity that cannot be formulated in mathematical or logical languages (Sugden 2009, p. 4).

Rubinstein (2006) likened economic models to fables or fairy tales. As such they are in his view only remotely related to reality, reaching absurd conclusions, and not directly testable. Grune-Yanoff and Schweinzer (2008) highlighted the role of stories in applying game theory. Cartwright (in press) compares models to parables and to fables, arguing that the latter have a moral, corresponding to a model's conclusion. Walliser (2011) provides an extensive taxonomy of the use of models in economics, ranging from the more standard scientific practices to those that are more specific to economics.

Viewing economic models as explorations, isolations, capacity-identifiers, and credible worlds are naturally not mutually exclusive. A special issue of *Erkenntnis* was devoted to this topic in 2009, focussing on the debate between the view of models as isolating tools and as credible worlds. While

Sugden (2009) insists on credible worlds as differing from isolating tools, Maki (2009) finds that they need not be so different, and Cartwright (2009) argues that they need to identify capacities in order to be successful (see also Knuuttila, 2009, and Kuorikoski and Lehtinen, 2009). Grune-Yanoff (2009) discusses the credibility of models, and how one can learn from minimal models, whereas Donato-Rodrıguez and Zamora-Bonilla (2009) view models as "inferential prostheses".

We share many of these views. Moreover, our main motivation, as well as our resolution to the sociological puzzles, have much in common with the works cited above. In particular, we agree that (i) economic models are often viewed differently than models in the other sciences; (ii) economic theory seems to value generality and simplicity at the cost of accuracy; (iii) models are expected to convey a message much more than to describe a well-defined reality; (iv) these models are often akin to observations, or to *gedankenexperiments*; and (v) the economic theorist is typically not required to clearly specify where her model might be applicable and how. Most importantly, we agree that economic models are used to suggest conclusions about real situations by means of inductive inference.

In this paper we attempt to take this view one step further, focusing on case-to-case induction, or on *case-based reasoning*, as a model of the way economists expect (some of) their models to be used. We propose a formal model of analogical reasoning, which can add to the understanding of some phenomena relating to the practice of economics.

3 Case-Based Scientific Reasoning

3.1 Case-Based and Rule-Based Reasoning

In everyday as well as professional life, people use both rule-based reasoning and case-based reasoning for making predictions, classifications, diagnostics, and for making ethical and legal judgments. Rule-based reasoning, in which the reasoner formulates general rules or theories, was formally introduced by the ancient Greeks, in the development of logic. In case-based (or, equivalently, analogical) reasoning, in contrast, the reasoner identifies similar past cases and uses those cases to guide the prediction (or classification, diagnosis, or ethical or legal judgement) in the current case. The term "case-based reasoning" was coined by Schank (1986) (see also Riesbeck and Schank (1989)). However, the discussion of this type of reasoning dates back to Hume (1748) at the latest.

Rule-based reasoning has several advantages over case-based reasoning. First, a rule is a concise description of a regularity, compared with a large and ever-growing database of cases that conform to this regularity. Second, formulating a small set of general rules gives people a feeling of understanding and explaining a phenomenon in a way that a database of cases does not. Thus, even if the two methods perform equally well in terms of prediction, there is a preference for rule-based approaches, and one is often willing to sacrifice some accuracy of prediction in return for the compactness of rules and the associated feeling of "cutting nature at its joints". However, when simple rules do not seem to be satisfactorily accurate, people might resort to case-based reasoning, making predictions in each problem by redrawing analogies to past cases in the database.

These two modes of reasoning exist also in statistics. Rule-based reasoning is akin to learning a distribution function, whereas case-based reasoning is related to data-based methods such as kernel estimation and nearest-neighbor approaches. However, the philosophy of science tends to view scientific activity as generating knowledge in the form of rules only. We argue that some of the practices that evolved in economics can be better understood if scientific knowledge can also be viewed as a collection of cases.

3.2 Economics as Case-Based Reasoning

3.2.1 How Does it Work?

We suggest that economic reasoning is partly case-based, and that one role of theory is to enrich the set of cases. That is, the analysis of a theoretical model can be viewed as an "observation" of a new case. Such a case is not real, but is a gedankenexperiment, an observation that is arrived at by pure logic. An observation of this type is new only to the extent that one has not thought about it before. But if the question has not been previously raised, or if the proof is not trivial, one learns something new by reading the result. (See Maki (2005) for a related view of models as experiments.)

Consider the following example. Akerlof's (1970) celebrated "lemons market" paper presents an example of buyers and sellers of used cars. The example makes certain general assumptions about the agents' behavior and information, as well as more specific assumptions and even particular numerical values. Under some such assumptions, it can be shown that the market will collapse completely. This example does not inform us of a new observation from the field or about a laboratory experiment. Nor is it a new finding from a long-forgotten archive or the result of a computer simulation. It is a mathematical proof, which happens to be rather obvious post-hoc. And yet, it is highly insightful, and economists tend to think that it has changed the way they think about markets.[3]

Despite the fact that this example can be stated as a mathematical result, it may be more useful to think about it as a case rather than as a general rule. As stated, the example can be viewed as the claim, "I have observed a case in which idealized agents, maximizing expected utility, with the following utility functions and the following information structure, behaved in such and such a way". The relevance of this observation for prediction will depend on the perceived similarity between the idealized agents and the real agents one is

[3]Indeed, Akerlof received the Nobel Prize in 2001 for this contribution.

concerned with, the similarity between the situation of the former and that of the latter, and so forth. An economist who is interested in real agents would therefore have to judge to what extent the situation she studies resembles the idealized situation in the "case" reported by Akerlof. (See Sugden, 2009, who argues that it is the reader who must make the mapping between a model and the reality it models.)

It is natural to think of experimental and empirical data as inputs for case-based reasoning as well. Indeed, the notion of *external validity* of an experiment involves the degree of similarity between the experiment and the real problem one is interested in. An economist who is asked to make a prediction in a given problem will then use case-based reasoning to learn from empirical data, experiments, theoretical models, and perhaps also historical examples, casual observations, and computer simulations. All cases, real, experimental, and theoretical, are aggregated, weighing their similarity and relevance, to generate predictions for the case at hand. In this sense case-based reasoning does not endow any type of information—empirical, experimental, or theoretical—with any privileged status.

When one engages in rule-based reasoning, one is expected to state rules that are accurate. To this end, the domain of applicability of the rules should be clearly defined. Observing counter-examples to the rule suggests that the rule has to be revised, or that its domain should be restricted. By contrast, when one employs case-based reasoning, there is no domain of applicability, and no universal statements are involved. A specification of the domain of applicability is replaced by a similarity judgment. This similarity judgment is often hinted at by the economist analyzing the model, but it is not part of the formal model. Moreover, the readers of a model may not agree with its author about its similarity to various problems. Rule-based knowledge is not complete without the "user's manual" that specifies the domain of applicability. By contrast, case-based knowledge allows for greater flexibility, separating the "hard" knowledge of cases from the "soft" judgment of similarity.

Rules can be refuted by cases.[4] By contrast, cases are not contradicted by other cases. Typically, for a given prediction problem different cases will suggest different predictions. The reasoner should then consider the totality of cases that make a certain prediction, judge their similarity, and compare it to that of each other possible prediction. The method applies even when some of the cases are theoretical. For example, assume that a theoretical analysis of the "ultimatum game" (Guth, Schmittberger, and Schwarze (1982)), in which utilities are only defined by monetary payoffs, suggests that player I will offer a minimal amount to player II, and that player II will accept the offer. Next assume that an experiment reveals a different outcome. If one conceives of the model as a general rule, one would have to conclude that the rule was violated, and perhaps re-define its scope of applicability. By contrast, if the theoretical analysis is construed as a case, as is the experimental result, the two coexist peacefully. Given a new prediction problem, an economist who is asked to make a prediction would have to ask herself, "is this real problem more similar to the theoretical analysis, assuming common knowledge of rationality with purely monetary payoffs, or is it more similar to the experiment?" In making this judgment the economist may draw on her knowledge of the players, the amounts of money involved, the time they have to make a decision, and so forth. Neither the theorist nor the experimentalist is expected to state *a priori* which real life problems belong to the same category as their case. Their job is only to contribute these cases as additions to the literature, and to leave similarity judgments to the practitioners who might use these cases in real life problems.

[4] Often, a rule is stated or interpreted probabilistically, and it can only be refuted statistically, that is, by a database of cases.

3.2.2 Is This Science?

Can case-based reasoning be a basis for science? The answer of course depends on the definition of "science", but it is useful to note that case-based reasoning can generate refutable claims if it is coupled with (i) an algorithm for the computation of similarity judgments and (ii) an algorithm for the generation of predictions based on judgments, such as kernel classification, a nearest-neighbor method, and the like. Should one commit to a similarity function and to the way in which it should be used, one would make predictions that can be tested and possibly refuted.[5]

The common practice in economic theory is to use models without a clear specification of the similarity function that should be used to apply them to concrete problems. (Again, see Sugden, 2009, who points out this phenomenon.) An economic theorist who offers a model prepares the ground for a practitioner who should employ her judgment in using this model; but the theorist's contribution falls short of a testable prediction.

Rule-based knowledge can also be suggested without the "user's manual" specifying the domain of applicability required to make the rule scientific. Indeed, proverbs may be viewed as universal statements that are made without a specification of the ranges of the variables over which one quantifies. Hence, in principle both rule-based and case-based knowledge can be presented without a specification of the way they should be applied. However, the empirical claim is that in the sciences rule-based knowledge tends to appear in a well-specified guise, whereas case-based knowledge in economics often does not.

[5]For example, if one uses kernel estimation, one may test hypotheses about the kernel (or similarity) function, as in Gilboa, Lieberman, and Schmeidler (2006). Kernel estimation is hardly a candidate for learning from theoretical cases, because repeatedly "observing" the same theoretical case does not add to our belief in its prediction. However, this example illustrates the general point, namely, that once one commits to a particular way in which the similarity function is to be applied, hypotheses about the similarity function become scientifically meaningful.

3.3 Revisiting the Puzzles

We argue that viewing economists as generating knowledge that is partly case-based explains the puzzles raised in Section 2. First, one need not wonder why economists feel that they gain insights and understand economics better using models whose assumptions are wrong. In the case-based approach, models cannot be wrong. As long as the mathematical analysis is correct, a theoretical case is valid, the same way that an empirical or experimental case is valid as long as it is reported honestly and accurately. Cases do not make any claim to generality, and therefore they cannot be wrong.

Consider the example of the ultimatum game. In the standard, rule-based model of science, the ultimatum experiment is a refutation of a rule, which should make one reject (or at least refine) the rule. But in the case-based model of science, the ultimatum experiment is but a case, as is the formal model, and economists should weigh both, along other cases, in making their predictions. Whether a case originates from empirical data, experiments, or theoretical analysis, it has the same epistemological stature for the economist.

This approach can also explain the high degree of mathematical sophistication in economics. One role of mathematical analysis is to obtain more observations, namely, theoretical cases. Similarly, analysis can extend the scope of existing cases. For example, if there is a proof that a certain result holds for two agents, and one proves that it holds for any number of agents, the new theoretical case may have a higher weight in further reasoning because it is more similar to some real cases of interest. In this sense, generalizing a mathematical result plays much the same role as repeating an experiment with participants drawn from a different population.

Finally, using the case-based view, one can also understand why economists and psychologists view their models differently. True to the standard, rule-based model of science, psychologists try to avoid refutations by being very explicit about the domain of applicability of their models. Economists, on the other hand, often offer models that are merely theoretical cases. These models cannot be refuted, and hence there is nothing to be lost by trying to draw analogies between them and new, remotely connected problems. On the contrary, every problem that may end up being similar to the model increases the model's popularity. As a result, economists have an incentive to view more real life cases as examples of their models, without risking their theory's reputation in so doing.

4 Related Phenomena

In this subsection we argue that the conceptualization of economic models as theoretical cases can also explain additional phenomena in the sociology of economics. The phenomena discussed here differ from the "puzzles" of Section 2 in that they are less conspicuous to academics outside economics.

4.1 Intuitiveness

Economists are often expected to provide intuition for their results, and it can be problematic for a result or a proof to be judged counterintuitive. As in the case of mathematics or theoretical physics, economic theory definitely values results that are difficult to prove. Indeed, in all of these disciplines results that are considered too obvious will typically not be published. However, in mathematics and in physics, once a non-trivial result has been established, one can hardly dismiss it based on its proof being counterintuitive. In contrast, in economics it appears to be legitimate for a referee to say "The proof is difficult, but, because I do not understand its intuition, I cannot support publication".[6] Why does economic theory value intuitive proofs? Why isn't

[6]We focus here on results that are supposed to have concrete implications in terms of economic behavior. There are mathematical results, such as Mertens and Zamir's (1985) formulation of the sense in which Harsanyi's type-space approach to incomplete information sacrifices no generality, that do not make any predictions in specific economic situations. These results are part of the theorists' discourse, used to convince economists that they should be using particular models and tools of analysis rather than others. Such results may be powerful rhetorical devices without their proofs necessarily being intuitive.

it sufficient for a result to be mathematically correct?

It might be necessary to first define what it means to say that an argument is intuitive. We suggest that an argument is judged to be intuitive if the various steps of the argument bear similarities to existing cases. For example, Newtonian physics is relatively intuitive because we are acquainted with billiard balls, and an argument that particles behave as they do because they have much the same properties as do billiard balls makes the scientific explanation familiar. By contrast, the quantum mechanics view of particles is less intuitive because the supporting arguments do not bring to mind any familiar concepts from our everyday experiences. Along similar lines, thinking of the relationship between a nucleus of an atom and the electrons as the relation between the sun and the planets is intuitive because it reminds us of phenomena we already know.[7] Thus, an argument is more intuitive, other things being equal, the more cases it reminds us of, and the stronger is the association (or, the greater the similarity) between the steps of the argument and these cases.

With this view of intuitiveness, let us consider an economic model as a theoretical case. Having a prediction problem at hand, the reasoner needs to compare the case to that problem, and judge their similarity, which will determine the relevance of the case to the prediction problem. However, the case-based view of economics does not restrict the similarity judgment to the assumptions of the model; in fact, the judgment is often performed for an entire proof, as if it were a story. Furthermore, each step in the proof may bring to mind other analogies, between the prediction problem and real past cases.

For example, consider the relevance of Akerlof's model to a given prediction problem. Judging the similarity of the model to the problem, one should ask, how similar are the agents in the model to the agents in reality? Are

[7]This analogy is nowadays considered misleading. Thus, modern physics can be said to view the similarity between the two systems as superficial.

the people in the real problem expected utility maximizers like the players in the model? Do the former entertain subjective probabilities as do the latter? And so on. But one can also look at the first step in the proof, and ask whether the result of that step is familiar from other cases. For instance, if the proof begins by suggesting that buyers will realize that they face a product of uncertain quality, and therefore might not be willing to pay too high a price for it, the reader might well be reminded of real cases in which quality was an unobserved variable, resulting in a lower price of the good. The fact that this step in the proof brings to mind real past cases, and that these make certain predictions more vivid, help to convince the reader that the theoretical case is relevant to the problem at hand.

We do not claim that the preference for intuitiveness is a clear-cut proof that economic models are perceived as cases rather than as rules. Indeed, one may attempt to make an argument for intuitiveness also in a rule-based view of science, arguing that our degree of belief in general assumptions is bolstered by similarity to known instances. Yet, if one subscribes to the classical view of science, according to which one relies on empirically valid assumptions and derives conclusions from them, one should not be allowed to rule out theoretical results based on the absence of an intuitive explanation of their proofs. Thus, we find the high value placed on intuitiveness as supporting the case-based view of economic models more than the rule-based one.

4.2 Axiomatizations

Economic theory seems to value axiomatic derivations of models of individual decision making, even when the models and their implications are well-known. For example, Rozen [35] provides an axiomatic derivation of intrinsic habit formation models that have appeared in the literature. Maccheroni, Marinacci, and Rustichini [28] axiomatized the general class of "variational preferences" and Strzalecki [40] axiomatized the class of "multiplier preferences" used by Hansen and Sargent [22]. Again, these axiomatizations were done long after the decision rules had been incorporated into economic theories. One may therefore ask, why does the profession value the exploration of foundations when a theory is already developed? Shouldn't the theory be directly tested based on its predictions, their fit to reality, and so forth?[8]

While there are many reasons to be interested in axiomatic derivations of behavioral models, we hold that the case-based view of economic theory explains the interest in axiomatizations better than does the rule-based view. Consider a simple, textbook example. Economists typically assume that each agent maximizes a utility function. This assumption is supported by an axiomatic derivation, saying that a preference relation that satisfies basic requirements of completeness and transitivity can (in a finite set-up) be represented by maximization of a certain function.

Such an axiomatic derivation is a characterization theorem. As such, it cannot make a theory more or less accurate. If we were to test how many economic agents do indeed maximize a utility function, or how many have a preference relation that is complete and transitive, we would necessarily obtain the same results, and conclude that the theory has the same degree of accuracy in its two equivalent representations. Moreover, when statistical errors are taken into account, one may argue that it is better to test the theory directly, rather than to separately test several conditions that are jointly equivalent to the theory. Hence, if economists were taking their theories as general rules that should fit the data, axiomatizations would be of little value for the selection of theories.

Now consider the case-based view of economic theory. According to this view, no general claim is made about economic agents. Rather, the economic theorist suggests certain theoretical cases in which agents who maximize a utility function behave in certain ways. These theoretical cases are to

[8]We are not dealing with a marginal phenomenon. All three axiomatizations quoted here were published in the best theory journal.

be judged according to their similarity to real prediction problems. When we ask ourselves, "Are people in this problem similar to the agents in the model?", we may indeed find out that different representations of the same mathematical structure result in different similarity judgments. For example, one might find it unlikely that a randomly chosen consumer would consciously maximize a utility function, but, at the same time, quite plausible that the consumer's decisions would respect transitivity. Thus, axiomatizations (in this case, of utility maximization) point out to us similarities that are not obvious *a priori*.[9]

In other words, we argue that the field values axiomatic derivations because axiomatizations and, more generally, equivalence theorems, can be powerful rhetorical tools. The standard view of science leaves little room for rhetoric: theories are confronted with the data, and should be tested for accuracy. By contrast, the case-based view of science lets rhetoric occupy center stage: scientists only offer cases, and these should be brought to bear upon prediction problems, where similarity and relevance should be debated as in a court of law.[10] With this openly-rhetorical view of science, the importance of axiomatizations is hardly a mystery.

5 A Formal Model

In this section we provide a formal model of analogical reasoning and rule based reasoning.[11] We then invoke some simple complexity results to provide insight into why economic reasoning often relies on analogies rather than rules, and why there is a powerful premium on simplicity in these analogies. We interpret our analysis in terms of economic problems, but there is no such

[9]Dekel and Lipman (2010) provide a similar motivation for axiomatic representations.

[10]McCloskey (1985) uses a more expansive notion of "rhetoric," encompassing means of persuasion that go well beyond debates over the relevant similarity function.

[11]We use the terms case-based reasoning and analogical reasoning interchangeably, preferring the former when we focus on the database of cases and preferring the latter when we focus on the analogies that give rise to similarity judgments.

formal restriction.

The following is an example of the kind of analogy we have in mind. Consider Spence's (1973) signaling model in which a worker chooses a level of education that signals her ability to potential employers. Suppose a student who has been taught this model is told that a new lawyer has come into town and has taken a five-year lease on his office, even though a six-month lease was available at the same rate. The student is asked why it might be optimal for him to have done this. We expect the student to see the analogy in which the worker is mapped into the lawyer, the firms are mapped into potential customers, high and low abilities for workers are mapped into high or low abilities for the lawyer, and the choice of education level is mapped into the choice of lease length. In Spence's model, firms make the inference that only high ability workers would find it profitable to choose a high level of education, while in the target problem only a high ability lawyer will be able to earn sufficient income to warrant a long-term lease.

An even more remarkable analogy would extend beyond hiring choices, and, indeed, beyond economics: a student of economic theory might be faced with an examination, asking why would a peacock "find" it evolutionarily advantageous to invest in carrying a heavy useless tail. Having seen Spence's model, one could expect the student to think in terms of signaling and rediscover Zahavi's famous "handicap principle" (Zahavi, 1975). While this example is clearly outside the realm of economics, it provides a useful test of students' understanding of the relevant economic principles.

The original Spence model of education choice is highly stylized and highly abstract. One could argue that we could make the model more realistic by including the possibility that some people go to college out of boredom and some for recreational reasons, by noting that there are different qualities of schools (two year colleges, state schools, elite colleges, ivy league universities), by allowing that students receive scholarships, and so on. Including such things will make the model more realistic, but carries two disadvantages:

first, the model may be too messy to admit any conclusions. As in any modeling activity, one faces a trade-off between accuracy and simplicity, and one may settle for a simpler model that can be analyzed and used for predictions rather than insist on a more accurate model that is too complex to analyze. Second, a more elaborate model of the education choice problem will make it much more difficult to see the analogy between that problem and, say, the lawyer problem discussed above (let alone the peacock problem). This reason to favor simplicity appears to be characteristic of case-based reasoning with an unspecified similarity function: because the readers of the model are supposed to seek analogies on their own, simplicity makes their task significantly easier.

Our model of analogical reasoning formalizes this latter property, namely that finding useful analogies between problems very quickly becomes extremely difficult when we make a model more realistic. In keeping with our view of models, our goal is to construct the simplest model of economic reasoning capable of making our points. We construct a model of economic reasoning as a prediction problem and then make our points in reverse order, examining the role of simplicity in effective analogical reasoning and then the trade-off between analogies and rules.

5.1 Prediction Problems

Our formulation of a prediction problem begins with a non-empty, finite set E. We interpret E as a set of objects that are the subject of analysis. This set will typically include people who act in real economic situations, agents in an economic model, or agents in an experiments. For example, in a real-life problem, E may include a person, say John Smith, and his choice of whether to go to college. In a model, E might include an agent, "Player I", and a selection of an education level from a binary set. In another real-life problem, E might include a lawyer as well as a choice of the length of his lease. In yet another problem, E might include a peacock rather than a person. The set might also include the preferences of these agents, acts that they choose, goods that they trade, and so on.

The essence of the problem, or its story, is told by a nonempty, finite set of *predicates* F. Intuitively, predicates identify the information available about the objects in E. Formally, a predicate is a function that maps tuples of elements in E into $\{0, 1\}$. For example, F might include a "1-place" predicate f that identifies, for every element of $e \in E$, whether that element is an agent ($f(e) = 1$). In the examples of the preceding paragraph, this predicate will identify John Smith, Player I, the lawyer, and the peacock as agents. Another 1-place predicate might capture behavioral assumptions, such as "this agent never chooses dominated strategies". A 2-place predicate might identify investment, so that we would have $f(e_i, e_j) = 1$ (for $e_i, e_j \in E$) if e_i invests in e_j. Presumably, this can only hold when e_i is an agent and e_j is a form of capital, such as higher education or a lease. As above, this investment relation may occur in a real-life example, in a laboratory experiment, or in a theoretical model. Similarly, 3-place predicates might describe preferences. For example, we might have $f(e_i, e_j, e_k) = 1$ if and only if e_i (an agent) prefers e_j (a consumption bundle) to e_k (another consumption bundle). In general, we interpret $f(e_1, \ldots, e_k) = 1$ as an indication that the property described by f is true of $(e_1, \ldots, e_k)$, and interpret $f(e_1, \ldots, e_k) = 0$ as an indication that the statement is either false or meaningless (without distinguishing these two possibilities). For example, a 2-predicate $f(e_i, e_j)$ might identify the location of (an agent) e_i as e_j. The value $f(e_i, e_j) = 1$ means that the location of e_i is indeed e_j (say, if e_i is "The President of the United States" and e_j is "Washington, D. C."), whereas $f(e_i, e_j) = 0$ might mean either that the statement is false (say, if e_i is "The President of the United States" and e_j is "Peoria, Illinois") or meaningless (say, if e_i is "Washington, D. C." and e_j is "The President of the United States").

A prediction problem is a pair (E, F). To consider analogies between prediction problems, we would like to start with two such problems, (E, F) and

(E', F'), and to consider mappings between them. Such mappings would map objects in E to objects in E' ("first-order analogies"), and also map predicates in F to predicates in F' ("second-order analogies"). To simplify the exposition, we restrict the argument in this section to first-order analogies, informally discussing second-order analogies in Section 6. It only reinforces our argument to note that our complexity result holds despite this significant simplification.

If we are to focus on first-order analogies, we need some way of saying that the predicates F (defined on E) are the same as the predicates F' (defined on E'). Toward this end, let $\mathcal{E}$ be the (finite) union of all sets of objects E that are within the purview of analysis, and consider predicates that are defined on $\mathcal{E}$, with the property that $f(e_1, \ldots, e_k) = 0$ if the objects $e_1, \ldots, e_k$ are *not* contained in a single set of objects E.[12] We will then simplify the notation by referring to prediction problems as simply E and E', understanding that we have a common set of predicates F that is defined on both problems (as well as on all of $\mathcal{E}$).

Thinking of the predicates in F being defined on all of $\mathcal{E}$ requires that these predicates be expressed in a sufficiently flexible language. Thus, instead of using the predicate "invests in higher education" for the Spence model and "invests in a long lease" for the lawyer example (or "invests in a heavy tail" for Zahavi's handicap principle), we assume that the relevant predicate $f \in F$ is simply "invests in some form of capital". Clearly, this very language already assumes a certain degree of abstraction that in turn simplifies the task of finding appropriate analogies. Once again, we note that our complexity result holds despite this simplification.[13]

[12]Hence, it is meaningful to say that a worker in a Spence signaling problem invests in education, and a peacock invests in a long tail, but not that the worker invests in a long tail.

[13]We interpret the set $\mathcal{E}$ as the set of all objects that might be the target of economic analysis. We cannot literally define $\mathcal{E}$ as the set of all objects without running into the classic paradoxes of set theory. We can instead think of $\mathcal{E}$ as containing the objects economists customarily examine. Different economists at different times may then have

The task of the analyst is to associate an outcome r with the prediction problem E. For simplicity we assume that outcomes are binary, that is, $r \in \{0, 1\}$. For example, the outcome might be whether trade occurs. This assumption is a simplification in two ways. First, an outcome can often be a real variable, or a vector of real variables, such as the level of inflation, level of employment and so forth. Second, it is implicitly assumed that the entire analysis focuses on a single question, so that the meanings of "0" and "1" are implicitly understood. In reality scientists collect data, run experiments, and analyze models that can be used for many different research questions, some of which may not even be specified at the time cases are collected. A more general model might describe outcomes as abstract entities, and capture their relevant aspects by functions that are defined on them (similar to the way predicates describe the prediction problem).

5.2 Analogies

One approach to prediction problems is to rely on analogies.

5.2.1 Cases

A *case* c is a prediction problem E coupled with its *outcome* r. If a case designates a data point that was empirically observed, the the prediction problem (including the values of its predicates) and r are observed simultaneously. In these cases, the economist can choose which entities and predicates to observe, but she typically cannot control the values of the predicates. For example, the economist might choose to observe whether trade takes place between individuals, and she can choose to focus on their endowments and preferences, but she has no control over the values of these variables. By contrast, if a case is an experimental observation, the experimenter is free

different views as to the scope of $\mathcal{E}$. Decisions to commit crimes or have children were once outside of economic analysis, but now are familiar, while one can still find differences of opinion as to whether $\mathcal{E}$ should include sets containing objects describing neural activity.

to set the values of the predicates, and the unknown is the outcome r. For example, an economist can decide to run an experiment in which she controls the participants' endowments and opportunities to trade, and observes whether they end up trading. Similarly, if the case is a theoretical study, the economist is free to assume any values of the predicates, and the outcome r is determined by mathematical analysis.

A *memory* is a finite collection of cases, M. The scientific challenge is to consider a memory M, and make a prediction about the outcome of a new prediction problem E.

5.2.2 Analogies—A formal definition

An *analogy* between prediction problem E and prediction problem E' is a 1-1 function $\varphi : E \to E'$. Prediction problem E will be referred to as the *origin* of the analogy, and prediction problem E' as its *target*.

The strength of the analogy depends on the values taken by the predicates over the set of objects. The analogy φ between E and E' will be considered a *perfect analogy* if, for every k- predicate f in F,

$$f(e_1, ..., e_k) = f(\varphi(e_1), ..., \varphi(e_k))$$

for every $e_1, ..., e_k \in E^k$. Thus, an analogy is perfect if all that is known about the prediction problems is identical.

The essence of analogical reasoning is to identify the similarity between each of the cases in the memory M and the prediction problem E, and then to make a prediction for E that is a function of the predictions and similarities in the memory. There are many details to be considered concerning the specification of this function, but we need only be concerned with the beginning step of this process—assessing the similarity between two prediction problems E and E'.

5.2.3 Finding Analogies

As was noted in Aragones, Gilboa, Postlewaite, and Schmeidler (2001), finding analogies is not a simple computational task. Even if one restricts attention to only two prediction problems and is interested only in the question of whether there exists a perfect analogy, the fact that the number of possible analogies grows exponentially in the number of predicates renders the problem intractable.

To make this precise, we borrow the notion of NP-Completeness from computer science.[14] A yes/no problem is NP if it is "hard" to find a solution, in the sense that there does not exist an algorithm that can find a solution in polynomial time (though it is "easy", in the sense that the task can be performed in polynomial worst-case time, to verify that a suggested solution is indeed a solution to the problem). NP-Completeness means more than this: for NP-Complete problems, if a polynomial algorithm can be found for one of them, it can be translated into polynomial algorithms for all other NP-Complete problems. Thus, a problem that is NP-Complete is at least as hard as many problems that have been extensively studied for years. We show:

Proposition 1 *The following problem is NP-Complete: given two prediction problems E and E', is there a perfect analogy $\varphi : E \to E'$ between them?*

To gain some intuition for this result, suppose that we have prediction problems E and E', with $|E| = k$ and $|E'| = n \geq k$. Then the number of 1-1 mappings $\varphi : E \to E'$ is

$$n\,(n-1)\ldots(n-k+1) = \frac{n!}{(n-k)!} = \binom{n}{k} k!$$

As k and n both increase, this number grows exponentially large (for example, it equals $n!$ when $k = n$). This does not necessarily imply that one cannot find

[14]See Appendix A for a short overview of the main concepts.

whether a perfect analogy exists in an efficient manner, but such exponential growth is perhaps the most obvious warning sign of computational difficulty.

Proof It is straightforward that the problem is in NP. To see that it is NP-Complete, observe that it is NP-Complete even if we restrict attention to $F = \{f\}$ and $f \in \mathcal{E}^2$. The analyst's task is then to determine whether, given two directed graphs, one is a sub-graph of the other. This problem is NP-Complete (for instance, the Clique problem can be reduced to it.) ■

It is easier to find analogies between prediction problems that do not have too many entities. In particular, suppose E is a theoretical model and consider the task of finding whether (and how) it applies to a prediction problem E'. As mentioned above, the set of all possible mappings from E to E' is of size

$$\frac{n!}{(n-k)!} \leq n^k$$

for $k = |E|$ and $n = |E'|$. For a fixed theoretical model E (and hence fixed k), however, the bound n^k is a polynomial in n. Hence, if k is sufficiently small, the computational task of finding analogies may be manageable, even if solved by brute force.

This leads to our first point: it is no surprise that economists prefer theoretical models with few "moving parts". A lower number of entities in the model makes it more likely that the model will be useful as a source of analogies for a prediction problem at hand.

5.3 Rules

Given a set of predicates F on the universal set of objects $\mathcal{E}$, a *rule* is formally defined as a prediction problem and an outcome, or (E, r), just as is a case. The distinction between rules and cases lies in the way they are used: a rule is interpreted as saying "*whenever* a set of entities E satisfies the relations defined by the predicates F, the result r will occur". For example, a rule may state that whenever there are two individuals who own one good each, and each prefers the good that the other has to her own, they will trade.

We emphasize that the mathematical object (E, r) can be used either as a case or as a general rule. In the preceding example, when $(E, 1)$ is interpreted as a case, we may think of it as saying, "once there were two individuals, i and j, who owned one good each, a and b respectively; each preferred the good owned by the other to her own, and they traded". Such a case could be an empirical observation or a result of an experiment. The case can also result from a theoretical analysis, if one adds to it appropriate assumptions such as "Agents i and j always reach Pareto efficient allocations".[15] However, none of these cases—empirical, experimental, or theoretical—is assumed to be a general theory, and thus none can be refuted by another case.

By contrast, when the case $(E, 1)$ is interpreted as a rule, a refutation requires only a single case to the contrary. In the example above, it suffices to have an experiment in which i and j are players, a and b are goods, i owns a and prefers b while j owns b and prefers a, and yet no trade occurs. More generally, a rule (E, r) is *refuted by* a case (E', r') if there is an analogy $\varphi : E \to E'$ between E and E' such that for every $f \in F$,

$$f(E) = f(E'),$$

but $r \neq r'$. That is, to determine that the case (E', r') refutes the rule (E, r) we first need to establish that the prediction problem E' indeed lies in the domain of applicability of the rule, given by the general template E. To this end, we need to verify that each of the predicates that hold in the statement of the rule also holds in the prediction problem. Only when it is established that the prediction problem is indeed an example of the general rule, will a different outcome $r' \neq r$ constitute a refutation of the latter.

Note that the definition of a refutation boils down to the definition of a perfect analogy. It follows that it is NP-Complete to determine, given a rule

[15] This assumption would have to be stated as a predicate, as would other behavioral assumptions about each agent separately or about several agents as a group.

(E, r) and a case (E', r'), does the case refute the rule.

We have thus shown that identifying the similarity of two cases is computationally difficult, as is identifying whether a case refutes a rule. In the case of analogical reasoning, our conclusion was that it would typically be expedient to work with simple cases, a finding that we interpreted as motivating the common penchant of economists for working with simple models. Similarly, we can expect economists to prefer simple rules, that is, generalizations (E, r) where $|E|$ is low.

Unfortunately (for rule-based reasoning), simple rules are easily refuted. Rule-based reasoning thus faces a challenging trade-off: complicated rules are computationally intractable, while simple rules are typically refuted. This leads to our second point: cases can never be refuted, and case-based reasoning is thus an attractive alternative to rule-based reasoning, allowing economists to work with models simple enough to be useful without worrying about refutations.

6 Standard Languages

6.1 Second-Order Analogies

Psychologists distinguish between different orders of analogies. First-order analogies are between objects for which the same predicates presumably hold. Second-order analogies are not only between objects, but also between the predicates. For example, comparing Mary's relationship with her advisor to John's relationship with his advisor is a first-order analogy. By contrast, comparing Mary's relationship with her advisor to John's relationship with his father is a second-order analogy, where the binary relation "is an advisor of" is likened to the binary relation "is a parent of".

Some of the more powerful and surprising analogies in economics are of second order. Consider, for example, Hotelling's (1929) famous model of two merchants on Main Street. The model predicts that the two will locate very close to each other, at the center of town (measured by the density of consumers along it). This is the equilibrium of the game played by the two sellers, assuming that the buyers choose to walk over to the seller who is closer to them. Indeed, any other location on the street by one seller allows the other seller to gain more than 50% of the market. As Hotelling notes, this model can be re-interpreted as a model of political competition, suggesting that two political candidates will express views that are centrist, for the same reasoning: assume that views are ordered on a line, and that every voter votes for the candidate whose expressed views are closest to the voter's. Under these assumptions, a candidate who expressed views that are not at the median allows her opponent to locate himself so that he gets more than 50% of the votes.

This analogy is particularly insightful because it is "cross-contextual": it relates different domains of knowledge. A priori the two stories are very different: one is about trade, the other about elections. In one story the key agents are trying to sell products and get a larger market share, whereas in the other they are politicians who attempt to draw votes. Indeed, the analogy is not perfect (in the informal sense) for these reasons: the merchants also determine prices, which do not have a clear equivalent in the political competition. Moreover, political candidates might have ideologies, or perceived ideologies, that restrict their freedom of location on the political opinion axis. Yet, the analogy certainly allows us to think about political competition in a new light, and to make some qualitative predictions that appear to be rather successful. Clearly, such an analogy is second-order: it not only maps voters to buyers, it also maps the predicate "votes for" to the predicate "buys from".

Consider another example. A principal-agent model might deal with a manager (the principal) who is trying to motivate workers (the agents) to exert effort even though their effort level is not directly observable. Such models have been analyzed extensively. Now compare this to a case in which

John insures his car. Should the car be damaged, the financial cost will be borne mostly by the insurance company, rather than by John himself. John might exert different levels of effort in trying to minimize the probability of such a damage, but his level of effort is not observable by the insurance company. Thus, the situation is akin to the principal-agent problem: one player (the worker, or John) can affect the expected payoff of another player (the principal, or the insurance company), where the latter cannot observe the action taken by the former. Principal-agent models of managers and workers are thus useful in understanding insurance markets.

This analogy is not immediately transparent. When John buys insurance, he is not employed by the insurance company. If anything, one would think of John as the customer who buys the insurance company services. Yet, when the possible acts and their outcomes are analyzed, it turns out that John is similar to the worker in affecting the other player's utility. This analogy is sometimes difficult to see, because the predicate "sells insurance to" in the insurance case is mapped to the predicate "hires" in the principal-agent case. Further, John, as the owner of the car, might be viewed as the more powerful principal, rather than as the agent whose services are hired. The analogy reverses the roles of buyer-seller, and yet it unveils a similar structure between two economic stories.

6.2 Standard Languages

Proposition 1 showed that finding first-order analogies is a daunting problem. Second-order analogies are yet more difficult to find, because they allow for a much richer set of possible mappings. When the analogical mapping only maps objects into objects, it is easier to search a database for possible analogies. Moreover, the words describing the predicates, such as "votes for", can serve as indices that allow one to search one's memory for cases that are similar to the prediction problem one is faced with. By contrast, when the analogical mapping allows "votes for" to be mapped to "buys from", there are many more possible analogies, and, worse still, the lexical indices provided by words do not suffice to bring to mind all the relevant cases.

One way to facilitate the task of finding second-order analogies is to use a standard language. One may view a "paradigm" or a "conceptual framework"[16] as consisting of a language that is supposed to be able to describe a large set of cases, coupled with certain principles for prediction. For example, the game-theoretical paradigm in economics starts with the language of players, strategies, information sets, outcomes, beliefs, and utilities. This language is somewhat abstract, but it allows economists to see cross-contextual analogies more easily. Once one replaces terms such as "voters", "buyers", "candidates", and "sellers" with the more abstract "players", one sees the analogy between the two stories that fit Hotelling's model. Similarly, when ownership and employer-employee relations are stripped from the stories, it is easier to understand why buying insurance is akin to working for a principal. In other words, a standard language allows one to see more similarities without resorting to second-order analogies. A paradigm may thus be useful even if it produces no rules.

7 Conclusion

There are fields of science that use standard languages, and that can also formulate general rules in these languages. This is arguably true of physics, whose standard language involves no more than five forces, and which succeeds in formulating theories that are both general and accurate. Unfortunately, the social sciences don't seem to be able to achieve this type of success. There are, in principle, two main directions in which a field might proceed: it can sacrifice generality for accuracy or vice versa.

When sacrificing generality, one would attempt to formulate rules that are supposed to hold only in very specific and well-defined situations. This

[16] See Gilboa and Schmeidler (2001).

is largely the direction taken by experimental psychology. It is also the way that part of economics is conducted. For instance, consider the textbook rule saying that demand goes down as the price goes up. To make sure that this rule is reasonably accurate, one may specify the domain of application so as to rule out speculative assets, goods of uncertain quality, or conspicuous consumption goods. With these restrictions, the rule appears to be a good approximation of the data.[17]

The other possible direction is to give up accuracy and aspire for generality in return. In an extreme version of this approach, one gives up the claim to formulate a general theory, so that accuracy is not an issue, but aims to have a language that describes a wide range of phenomena and allows for higher order analogies. Thus, rule-based reasoning is discarded in favor of case-based reasoning, and, in return, the latter becomes very powerful. The claim we are trying to make is that this is the direction taken by much of microeconomic theory is the past few decades, using game theory as the standard model, and generating insightful analogies rather than accurate rules.

[17]Giffen goods are a counter-example used in class, but they are certainly rare.

8 Appendix: Computational Complexity

A problem can be thought of as a set of legitimate inputs, and a correspondence from it into a set of legitimate outputs. For instance, consider the problem "Given a graph, and two nodes in it, s and t, find a minimal path from s to t". An input would be a graph and with two distinguished nodes. These are assumed to be appropriately encoded into finite strings over a given alphabet. The corresponding encoding of a shortest path between the two nodes would be an appropriate output.

An algorithm is a method of solution that specifies what the solver should do at each stage. Church's thesis maintains that algorithms are those methods of solution that can be implemented by Turing machines. Church's thesis is neither a theorem nor a conjecture, because the term "algorithm" has no formal definition. In fact, Church's thesis may be viewed as defining an "algorithm" to be a Turing machine. It has been proved that Turing machines are equivalent, in terms of the algorithms they can implement, to various other computational models. In particular, a PASCAL program run on a modern computer with an infinite memory is also equivalent to a Turing machine and can therefore be viewed as a definition of an "algorithm".

It is convenient to restrict attention to YES/NO problems. Such problems are formally defined as subsets of the legitimate inputs, interpreted as the inputs for which the answer is YES. Many problems naturally define corresponding YES/NO problems. For instance, the previous problem may be represented as "Given a graph, two nodes in it s and t, and a number k, is there a path of length k between s and t in the graph?" It is usually the case that if one can solve all such YES/NO problems, one can solve the corresponding optimization problem. For example, an algorithm that can solve the YES/NO problem above for any given k can find the minimal k for which the answer is YES (it can also do so efficiently). Moreover, such an algorithm will typically also find a path that is no longer than the specified k.

Much of the literature on computational complexity focuses on time complexity: how many operations will an algorithm need to perform in order to obtain the solution and halt? It is customary to count input/output operations, as well as logical and algebraic operations as taking a single unit of time each. Taking into account the amount of time these operations actually take (for instance, the number of actual operations needed to add two numbers of, say, 10 digits) typically yields qualitatively similar results.

The literature focuses on asymptotic analysis: how does the number of operations grow with the size of the input? It is customary to conduct worst-case analyses, though attention is also given to average-case performance. Obviously, the latter requires some assumptions on the distribution of inputs, whereas worst-case analysis is free from distributional assumptions. Hence the complexity of an algorithm is generally defined as the order of magnitude of the number of operations it needs to perform, in the worst case, to obtain a solution, as a function of the input size. The complexity of a problem is the minimal complexity of an algorithm that solves it. Thus, a problem is polynomial if there exists an algorithm that always solves it correctly within a number of operations that is bounded by a polynomial of the input size. A problem is exponential if all the algorithms that solve it may require a number of operations that is exponential in the size of the input, and so forth.

Polynomial problems are generally considered relatively "easy", even though they may still be hard to solve in practice, especially if the degree of the polynomial is high. By contrast, exponential problems become intractable already for inputs of moderate sizes. To prove that a problem is polynomial, one typically points to a polynomial algorithm that solves it. Proving that a YES/NO problem is exponential, however, is a very hard task, because it is generally hard to show that there does *not* exist an algorithm that solves the problem in a number of steps that is, say, $O(n^{17})$ or even $O(2^{\sqrt{n}})$.

A non-deterministic Turing machine is a Turing machine that allows multiple transitions at each stage of the computation. It can be thought of as a parallel processing modern computer with an unbounded number of processors. It is assumed that these processors can work simultaneously, and, should one of them find a solution, the machine halts. Consider, for instance, the Hamiltonian path problem: given a graph, is there a path that visits each node precisely once? A straightforward algorithm for this problem would be exponential: given n nodes, one needs to check all the $n!$ permutations to see if any of them defines a path in the graph. A non-deterministic Turing machine can solve this problem in linear time. Roughly, one can imagine that $n!$ processors work on this problem in parallel, each checking a different permutation. Each processor will therefore need no more than $O(n)$ operations. In a sense, the difficulty of the Hamiltonian path problem arises from the multitude of possible solutions, and not from the inherent complexity of each of them.

The class **NP** is the class of all YES/NO problems that can be solved in **P**olynomial time by a **N**on-deterministic Turing machine. Equivalently, it can be defined as the class of YES/NO problems for which the validity of a suggested solution can be verified in polynomial time (by a regular, deterministic algorithm). The class of problems that can be solved in polynomial time (by a deterministic Turing machine) is denoted **P** and it is obviously a subset of NP. Whether P=NP is considered to be the most important open problem in computer science. While the common belief is that the answer is negative, there is no proof of this fact.

A problem A is **NP-Hard** if the following statement is true ("the conditional solution property"): if there were a polynomial algorithm for A, there would be a polynomial algorithm for any problem B in NP. There may be many ways in which such a conditional statement can be proved. For instance, one may show that using the polynomial algorithm for A a polynomial number of times would result in an algorithm for B. Alternatively, one may show a polynomial algorithm that translates an input for B to an input for A, in such a way that the B-answer on its input is YES iff so is the A-answer of its own input. In this case we say that B is reduced to A.

A problem is **NP-Complete** if it is in NP, and any other problem in NP can be reduced to it. It was shown that the SATISFIABILITY problem (whether a Boolean expression is not identically zero) is such a problem by a direct construction. That is, there exists an algorithm that accepts as input an NP problem B and input for that problem, z, and generates in polynomial time a Boolean expression that can be satisfied iff the B-answer on z is YES. With the help of one problem that is known to be NP-Complete (**NPC**), one may show that other problems, to which the NPC problem can be reduced, are also NPC. Indeed, it has been shown that many combinatorial problems are NPC.

NPC problem are NP-Hard, but the converse is false. First, NP-Hard problems need not be in NP themselves, and they may not be YES/NO problems. Second, NPC problems are also defined by a particular way in which the conditional solution property is proved, namely, by reduction.

There are by now hundreds of problems that are known to be NPC. Had we known one polynomial algorithm for one of them, we would have a polynomial algorithm for each problem in NP. As mentioned above, it is believed that no such algorithm exists.

One NPC problem is the "clique problem": given a graph, find the largest completely connected subgraphs.

References

[1] Hirotugu Akaike. An approximation to the density function. *Annals of the Institute of Statistical Mathematics*, 6(2): 127–132, 1954.

[2] George A. Akerlof. The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *Quartely Journal of Economics*, 84(3): 488–500, 1970.

[3] Enriqueta Aragones, Itzhak Gilboa, Andrew Postlewaite, and David Schmeidler. Rhetoric and Analogies. , mimeo., 2001.

[4] S. Brue. *The Evolution of Economic Thought, 5th edition.* Houghton Mifflin Harcourt, 1993.

[5] Nancy Cartwright. Capacities. *in J. Davis, W. Hands, and U. Maki (Eds.), The handbook of economic methodology* (pp. 45–48). Cheltenham: Edward Elgar, 1998.

[6] Nancy Cartwright. Models: Parables vs. Fables. *in Beyond Mimesis and Convention: Representation in Art and Science, Roman Frigg, Matthew Hunter eds.* In Press.

[7] Nancy Cartwright. If No Capacities Then No Credible World. But Can Models Reveal Capacities?. *Erkenntnis*, 70(1): 45–58, 2009.

[8] Eddie Dekel and Barton L. Lipman. How (Not) to Do Decision Theory. *Annual Review of Economics* 2:257–282, 2010.

[9] Xavier de Donato Rodriguez and Jesus Zamora Bonilla. Credibility, Idealisation, and Model Building: An Inferential Approach. *Erkenntnis*, 70(1): 101–118, 2009.

[10] Evelyn Fix and J. L. Hodges. Discriminatory Analysis. Nonparametric Discrimination: Consistency Properties. Technical report 4, project

number 21-49-004, USAF School of Aviation Medicine, Randolph Field, Texas, 1951.

[11] Evelyn Fix and J. L. Hodges. Discriminatory Analysis. Nonparametric Discrimination: Small Sample Performance. Report A193008, USAF School of Aviation Medicine, Randolph Field, Texas, 1952.

[12] Milton Friedman. The Methodology of Positive Economics. In *Essays in Positive Economics*, pp. 3–43, 1953.

[13] Allan Gibbard and Hal R. Varian. Economic Models. *The Journal of Philosophy*, 75(11): 664–677, 1978.

[14] Gabrielle Gayer and Itzhak Gilboa. Analogies and Theories: The Role of Simplicity and the Emergence of Norms, mimeo., 2012.

[15] Itzhak Gilboa, Offer Lieberman, and David Schmeidler. Empirical Similarity. *Review of Economics and Statistics*, 88: 433–444, 2006.

[16] Itzhak Gilboa, Larry Samuelson, and David Schmeidler. The Dynamics of Induction in a Unified Model, mimeo., 2012.

[17] Itzhak Gilboa, and David Schmeidler. *A Theory of Case-Based Decisions*. Cambridge University Press, Cambridge, 2001.

[18] Werner Guth, Rolf Schmittberger, and Bernd Schwarze An Experimental Analysis of Ultimatum Bargaining. *Journal of Economic Behavior and Organization* , 3(4): 367–388, 1982.

[19] Till Grune-Yanoff. Preface to 'Economic Models as Credible Worlds or as Isolating Tools?'. *Erkenntnis*, 70(1): 1–2, 2009.

[20] Till Grune-Yanoff. Learning from Minimal Economic Models. *Erkenntnis*, 70(1): 81–99, 2009.

[21] Till Grune-Yanoff and Paul Schweinzer The Roles of Stories in Applying Game Theory. *Journal of Economic Methodology* , 15: 131–146, 2008.

[22] Lars P. Hansen and Thomas J. Sargent. Robust Control and Model Uncertainty. *American Economic Review*, 91: 60-66, 2001.

[23] Daniel Hausman. *The Inexact and Separate Science of Economics*. Cambridge University Press, Cambridge, 1992.

[24] Harold Hotelling. Stability in Competition. *Economic Journal*, 39(153): 41–57, 1929.

[25] David Hume. *An Enquiry Concerning Human Understanding*. Clarendon Press, Oxford, 1748.

[26] Tarja Knuuttila. Isolating Representations Versus Creible Constructions? Economic Modelling in Theory and Practice. *Erkenntnis*, 70(1): 59–80, 2009.

[27] Jaakko Kuorikoski and Aki Lehtinen. Incredible Worlds, Credible Results. *Erkenntnis*, 70(1): 119–131, 2009.

[28] Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini. Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. *Econometrica* 74: 1447–1498, 2006.

[29] Uskali Maki. Isolation, Idealization and Truth in Economics. *Poznan Studies in the Philosophy of the Sciences and the Humanities* 38: 147–168, 1994.

[30] Uskali Maki. Models are Experiments, Experiments are Models. *Journal of Economic Methodology* 12: 303–315, 2005.

[31] Uskali Maki. MISSing the World. Models as Isolations and Credible Surrogate Systems. *Erkenntnis*, 70(1): 29–43, 2009.

[32] Deirdre McCloskey. *The Rhetoric of Economics*. University of Wisconsin Press, Madison, 1985.

[33] Jean–François Mertens and Shmuel Zamir. Formulation of Bayesian Analysis for Games with Incomplete Information. *International Journal of Game Theory* 14: 1–29, 1985.

[34] Christopher K. Riesbeck and Roger C. Schank. *Inside Case-Based Reasoning*. Lawrence Erlbaum Associates, Hilldale, New Jersey, 1989.

[35] Kareen Rozen. Foundations of Intrinsic Habit Formation, *Econometrica*, 78 (4): 1341–1373, 2010.

[36] Ariel Rubinstein. Dilemmas of an Economic Theorist. *Econometrica* 74: 865–883, 2006 (Presidential Address to the Econometric Society, Madrid 2004).

[37] Roger C. Schank. *Explanation Patterns: Understanding Mechanically and Creatively*. Lawrence Erlbaum Associates, Hilldale, New Jersey, 1986.

[38] Bernard W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, London and New York, 1986.

[39] Michael Spence. Job Market Signaling. *Quarterly Journal of Economics* 87 (3): 355–374, 1973.

[40] Tomasz Strzalecki. Axiomatic Foundations of Multiplier Preferences. *Econometrica* 79: 47–73, 2011.

[41] Robert Sugden. Credible Worlds: The Status of Theoretical Models in Economics. *Journal of Economic Methodology*, 7(1): 1–31, 2000.

[42] Robert Sugden. Credible Worlds, Capacities, and Mechanisms. *Erkenntnis*, 70(1): 3–27, 2009.

[43] Robert Sugden. Explanations in Search of Observations. *Biology and Philosophy*, 26(5): 717–736, 2011.

[44] Amos Tversky. Features of Similarity. *Psychological Review*, 84(4): 327–352, 1977.

[45] Bernard Walliser. *Comment raisonnent les economistes: les fonctions des modeles*. Odile Jacob, Paris, 2011.

[46] Amotz Zahavi. Mate Selection – A Selection for a Handicap. *Journal of Theoretical Biology*, 53 (1): 205–214, 1975.