

1 Introduction

The language of game theory—coalitions, payoffs, markets, votes—suggests that it is not a branch of abstract mathematics; that it is motivated by and related to the world around us; and that it should be able to tell us something about that world. Most of us have long realized that game theory and the “real world” (it might better be called the complex world) have a relationship that is not entirely comfortable; that it is not clear just what it is that we are trying to do when we build a game-theoretic model and then apply solution concepts to it. This is the subject I would like to explore in this paper. I might add that much the same questions apply to economic theory, at least the kind that those of us working in mathematical economics see most often; and that much of my paper will apply, *mutatis mutandis*, to economic theory as well. There is a branch of philosophy that deals with theory in the social sciences, so some of the things I have to say are unquestionably old hat. But I am not trying to be particularly original: I am only trying to open this topic, which I think concerns all of us, for discussion, and to suggest a particular point of view. No doubt others have thought about these questions more thoroughly and deeply than I have, and are better versed in the history and philosophy of science in general. I will be grateful to anybody who sets me straight when I err, and who gives me references for the things I get right.

My main thesis is that a solution concept should be judged more by what it does than by what it is; more by its success in establishing relationships and providing insights into the workings of the social processes to which it is applied than by considerations of *a priori* plausibility based on its definition alone.

The first eight sections of the paper are concerned with generalities; sections 9–17 illustrate my point of view; and the last section summarizes and concludes.

This chapter originally appeared in *Frontiers of Economics*, edited by K. Arrow and S. Honkapohja, Basil Blackwell, Oxford, 1985. Reprinted with permission.

This work was supported by National Science Foundation Grant SES-8201373 and Contract ONR-N00014-79-C-0685 at the Institute for Mathematical Studies in the Social Sciences (Economics), Stanford University, California. The paper presents an approach to the philosophy of Game Theory; it is not a survey. The examples and references are an eclectic, largely haphazard selection—strongly skewed towards applications—from a truly enormous literature. The idea was merely to illustrate some of the points in the text; the references are neither representative nor systematically chosen, even within the particular subjects actually discussed.

2 Comprehension

To come to grips with the question of what we are trying to do in game theory, we must first back off a little and ask ourselves what science in general is trying to do. The “man in the street” may answer this question in terms of practical applications: light bulbs, plastic, computers, atom bombs, preventing depression, and so on. He concedes that, in the long run, applications and inventions require a broad infrastructure of “basic” science; but (according to this view) the *purpose* of science is the development of the practical application.

More sophisticated observers, including many scientists themselves, answer the question in terms of predictive power. The theory of relativity was a success, they believe, because it predicted the precession of the perihelion of Mercury and the displacement of the images of the stars in a solar eclipse. If a theory has no predictive power—if it is not “falsifiable”—then it is not science.

Both these views miss the main point, I think. On the most basic level, what we are trying to do in science is to understand our world. Predictions are an excellent means of testing our comprehension, and once we have the comprehension, applications are inevitable; but the basic aim of scientific activity remains the comprehension itself.

3 Three Components of Comprehension: Relationships, Unification, Simplicity

Comprehension is a complex concept, with several components. Perhaps the most important component has to do with fitting things together, relating them to each other. To “understand” an idea or a phenomenon—or even something like a piece of music—is to relate it to familiar ideas or experiences, to fit it into a framework in which one feels “at home.” When one first listens to Bach, one feels attacked by a discordant jumble of meaningless, disconnected sounds. But eventually one begins to hear patterns; the horn takes up what the violin said, groups of sound recur in minor and major keys, passages are repeated. Things jell, and one begins to feel “at home.” After a while one recognizes the style, and even if one is listening to an unfamiliar piece one can relate it to other pieces by the same composer or of the same period. One understands the music.

I would like to emphasize that I am not talking merely about familiarity; while it is important, it is not the main point. I am talking about relating, associating, *recognizing patterns*. Snowflakes are hexagonal; the shells of certain snails are logarithmic spirals; buses on busy routes arrive

in bunches; in their orbits around the sun, the planets sweep out equal areas in equal times; waves and ripples occur in the ocean as well as on sand dunes; fever is associated with infections; E predominates as the Deity's name in some parts of the bible, J in others; even total randomness has its patterns (normal and Poisson distributions, etc.). After a while, a pattern becomes so familiar that it itself is viewed as an observation, and then we begin seeking—and finding—patterns in the patterns. The law of gravitation is a pattern that one can discern both in the motion of the planets and in the form of a water wave or a sand wave. And there are also different kinds of waves—electromagnetic or sound waves, which seem to have nothing to do with gravity. Crystalline form is observed not only in snow, but in many other materials. And so on.

This brings us to the second component of comprehension, which is really part of the first: unification. The broader the area that is covered by a theory, the greater is its “validity.” I am not thinking of “validity” in the usual sense of truth, but rather in the sense of applicability or usefulness; I am measuring the validity of an idea by the amount that people (directly or indirectly) use it. Part of the greatness of theories like gravitation or evolution, or the atomic theory of matter, is that they cover so much ground, that they “explain” so many different things.

Of course, a unificatory theory is really a special case of a relationship; many different phenomena are pulled together—related to each other—by means of it. The idea of gravitation in itself, in the abstract, is rather mysterious; it is important because it enables us to relate the tides to the motion of the planets and to the trajectories of shells and missiles.

The third component of comprehension that I would like to discuss is *simplicity*. What I mean is mostly the opposite of complexity, though the other meaning of “simple”—the opposite of “difficult”—also plays a role. Here there are several sub-components. One is sparseness; as few as possible exogenous parameters should be used to account for any particular phenomenon. Ptolemy's theory of epicycles is in a certain sense correct; it cannot be rejected on the basis of the evidence. But to describe the motion of a heavenly body one needs many exogenous parameters, whereas in Newton's theory one needs just two (mass and velocity).

In addition to sparseness in exogenous parameters, one would like sparseness in the basic structure of the theory. Here again Newton's theory is an apt example, because it can be derived from his three basic laws of motion plus the inverse square law of gravitation, all of which are very simple. The theory of evolution or the atomic theory of matter are other apt examples of sparseness, or relative sparseness, in basic structure. An example of complexity—the opposite of what we want—is modern elementary particle theory. It is clear that nobody is particularly happy

with this, and that it is considered an intermediate stage on the way to a more satisfactory theory.

We come finally to the matter of simplicity in the sense that is opposite to difficulty. For a theory to be useful, working with it must be practical. If you cannot figure out what it implies, it won't unify anything, it won't establish relationships. If the chief problems of celestial mechanics had been 3-body (or n -body) problems, then the theory of gravitation would not have been very important, because we could not have used it to calculate any results. All other things being equal, the simpler a theory is, the more useful it is—and hence the more valid.

4 Science and Truth

Most readers will by now have understood that, in my view, scientific theories are not to be considered “true” or “false.” In constructing such a theory, we are not trying to get at the truth, or even to approximate to it: rather, we are trying to organize our thoughts and observations in a useful manner.

One rough analogy is to a filing system in an office operation, or to some kind of complex computer program. We do not refer to such a system as being “true” or “untrue”; rather, we talk about whether it “works” or not, or, better yet, how well it works. As an office operation grows, filing systems change and evolve; at some points, completely new and different systems may be introduced, to accommodate to the evolution of the kind and amount of material to be filed.

Similarly, scientific theories must be judged by how well they enable us to organize and understand our observations; by how well they “work.” As our observations increase in volume and change in character, old scientific theories are no longer as appropriate as before; they either evolve and change, or are replaced by entirely new and different theories. Truth, however, is not the issue. We discard a theory not because it has been “disproved,” but because it no longer works, is no longer appropriate.

It is even possible for two competing theories to exist happily side by side and be used simultaneously, in much the same way that many of us file letters both chronologically and by the name of the correspondent. Two famous examples are relativistic vs Newtonian mechanics, and wave vs particle theories of light. In each case each of the theories has its areas of usefulness. I remember reading in my teens that one famous scientist considered the wave theory “true” on Mondays, Wednesdays and Fridays, whereas he preferred the particle theory on Tuesdays, Thursdays and Saturdays (apparently he didn't work on Sundays). He

used either, at convenience, on any day of the week; but the problem of which one was “true” seems to have bothered him. Apparently he didn’t subscribe to my view of science; or maybe his flippant remark was meant to indicate that he did, that the matter of “truth” was secondary to him, and that the most important thing was to get on with his work.

The example of relativistic vs Newtonian mechanics is particularly instructive. It is probably fair to say that most of those scientists who seek the “truth” consider relativistic mechanics a better approximation to the truth than Newtonian mechanics. Nevertheless, they continue to use Newtonian mechanics for most “everyday” purposes. Why? Well, they say, the Newtonian theory is usually a good enough approximation to relativity. Why settle for an approximation when you can get it exactly right? Well, they might say, in many cases the relativistic theory is too cumbersome to work with; the Newtonian theory is more workable. But then, it appears, “truth” is after all not the only criterion; Newtonian mechanics continues to be used as a model—on a much larger scale than relativity—even after it has been discredited from the point of view of “truth.”

I vividly remember an afternoon ten or fifteen years ago when Paul Erdos, on one of his many visits to Jerusalem, was delivering a lecture at the mathematics colloquium. Though his mathematical powers were (and are) still extraordinary, he started by saying that he was now old and exhausted (“Zaken Vetashush”), and that little could be expected of him. Of course we all protested. But Erdos insisted, and finally stated flatly that he could prove that he was no less than 2 billion years old. In his childhood the earth had been 2 billion years old; now it was 4 billion years old; the conclusion about his age was inescapable.

Like many good jokes, Erdos’s has a serious kernel. In the 1920s, the model that best fitted the known observations and existing theory was a 2-billion-year-old earth. By the 1970s, radioactive dating had been discovered, our ways of thought had changed in many ways, and the 4-billion-year model fitted much better. Evidence that had seemed strong and convincing 50 years ago gave way to stronger and more convincing evidence in a different direction; the older evidence had to be, and was, explained away. It would be foolhardy to think that the process has ended here, that no new evidence will be discovered, that the earth is *truly* 4 billion years old. It seems much more likely that in the course of time we will change our minds—or rather our model—again and again; indeed, the end of the process is not in sight. Why say that we were “wrong” then, that we have discovered the errors of our ways and are “right” now? It seems much more apt to say that when Erdos was a child, 2 billion years was right, and now 4 billion years is right. Each of

these two models is the one that best fits, or organizes, or ties together, the observations available at its time, and the theories current at its time.

Here is another example. When the red shift in light from distant galaxies was first observed, two theories presented themselves. The one was that, in travelling such long distances, the light somehow gets “tired,” the frequency of vibration decreases, and a shift to the red results; the other theory was that the galaxies are receding—the universe is expanding—and the red shift is associated with the Doppler effect. At present, the expanding universe theory appears to have won out. Most reasonable scientists will agree, however, that this area is very much in flux, and that our thinking on these matters is liable to change. Some will express this by saying that, though we now think that the universe *really is* expanding, new evidence—or new theories—may change our minds. But I prefer to say that the best model or system we can find for organizing our current observations and current theories involves an expanding universe; and that in the future we may, for one reason or another, find that other systems work better.

Some philosophies deny altogether the existence of objective truth, but for my purposes this is not necessary, and I do not wish to insist on it. The concept of truth applies to *observations*; one can say that such and such were truly the observations. It also applies to all kinds of everyday events, like whether or not one had hamburger for dinner yesterday. It does not, however, apply to *theories*.

I'd like to stress that I am not being dogmatic. What I would like to sell is a point of view, not a hard and fast position. The boundary lines are without doubt fuzzy. A conjecture about what happened next door yesterday may well be considered true or false, whereas a theory about what happened a billion years ago or a billion light-years away will not be. It can even happen that a theory graduates into a truth; whereas the roundness of the earth was a theory for the ancients, it can, I think, be considered a truth for us.

Readers may ask, why am I so insistent on making this point? What difference does it make whether we are looking for the truth or for a workable model, as long as we are not dogmatic and are willing to consider new evidence or new ways of thinking?

I think that the distinction is crucial for social science in general, and for game theory and economics in particular. For one thing, we have the matter of pluralism—the existence of parallel scientific theories side by side. We have seen that this occurs even in the natural sciences; but in our disciplines it is ubiquitous, it is very much the name of the game. People ask, since game theory offers a multiplicity of solution notions, what good can it be? Which solution notion is the right one? How do people

“truly” behave? If one takes the point of view suggested above, this question loses much of its sharpness. None of the solution notions tells us how people truly behave. They do not go about organizing blocking coalitions, as the core might suggest; they do not object and counter-object as in the bargaining set; they do not declare dividends as in Harsanyi’s value; and so on. Rather, a solution notion is the scientists’ way of organizing in a single framework many disparate phenomena and many disparate ideas. And it is simply that the kind of social phenomenon that our disciplines cover has not yet proved amenable to a single overriding system of thought, and perhaps never will.

Another reason for stressing this point—that a theory should not be thought of as true or false—is to avoid the pitfalls of taking it too literally. For example, an objection that has been raised to the fundamental notion of utility maximization is that, for one reason or another, individuals do not *really* maximize utility. Alternatives such as satisficing have been proposed, which sometimes seem more appropriate as descriptions of true individual behavior. But the validity of utility maximization does not depend on its being an accurate description of the behavior of individuals. Rather, it derives from its being the underlying postulate that pulls together most of economic theory; it is the major component of a certain way of thinking, with many important and familiar implications, which have been part of economics for decades and even centuries. Alternatives such as satisficing have proved next to useless in this respect. While attractive as hypotheses, there is little theory built on them; they pull together almost nothing; they have few interesting consequences. In judging utility maximization, we must ask not “Is it plausible?” but “What does it tie together, where does it lead?”

5 Game Theory as Descriptive Science

Briefly put, game and economic theory are concerned with the interactive behavior of *Homo rationalis*—rational man. *Homo rationalis* is the species that always acts both purposefully and logically, has well-defined goals, is motivated solely by the desire to approach these goals as closely as possible, and has the calculating ability required to do so.

The difficulty with this definition is apparent as soon as one writes it down. *Homo rationalis* is a mythical species, like the unicorn and the mermaid. His real-life cousin, *Homo sapiens*, is often guided by subconscious psychological drives, or even by conscious ones, that are totally irrational; herd instincts play a large role in his behavior; even when his goals are well-defined, which isn’t often, his motivation to achieve them may be less than complete; far from possessing infinite calculating ability,

he is often downright stupid; and even when intelligent, he may be tired or hungry or distracted or cross or drunk or stoned, unable to think under pressure, able to think only under pressure, or guided more by his emotions than his brains. And this is only a very partial list of departures from the rational paradigm.

Thus we cannot expect game and economic theory to be descriptive in the same sense that physics or astronomy are. Rationality is only one of several factors affecting human behavior; no theory based on this one factor alone can be expected to yield reliable predictions.

In fact, I find it somewhat surprising that our disciplines have any relation at all to real behavior. (I hope that most readers will agree that there is indeed such a relation, that we do gain *some* insight into the behavior of *Homo sapiens* by studying *Homo rationalis*.) There is apparently some kind of generalized invisible hand at work. While in any given situation an individual may well act irrationally, there seems to be a cumulative effect of numbers and time and learning that pushes people “in general” in the direction of rational decision-making. It doesn’t make people more and more rational, but as a given setting gets more and more common and familiar, it makes them act more and more rationally *in that setting*. The descriptive power of the rationality hypothesis is nicely summed up in Abraham Lincoln’s remark, “You can fool all of the people some of the time, and some of the people all of the time, but you can’t fool all of the people all of the time.” If one is careful not to expect too much, then *Homo rationalis* can serve as a *model* for certain aspects of the behavior of *Homo sapiens*. This is related to ideas in biology and evolution, in which the doctrine of survival of the fittest translates into maximizing behavior on the part of individual genes. We know that genes don’t “really” maximize anything; but the phenomena we observe, or some of them, are nicely tied together by the hypothesis that they act *as if* they were maximizing. Things are more complicated in the social sciences, first because the decisions themselves are very complex, and second because non-maximizing conduct is not as ruthlessly punished as in the jungle; but perhaps there is a similar trend.

Descriptively speaking, then, we can expect our disciplines only *sometimes* to explain or provide insights into “real” phenomena. We cannot expect them always to do so, because they are admittedly incomplete. We cannot even say beforehand *when* we expect them to do so, because we do not yet know how to integrate rational sciences like game theory and economics with non-rational sciences like psychology and sociology to yield accurate predictions. The criterion for judging our theories cannot be rigid; we cannot ask, is it right or is it wrong? Rather, we must ask, *how often* has it been useful? *How* useful has it been?

All this may sound very slippery and unsatisfactory. There are no firm predictions, no falsifiability. If our theory appears not to work, we don't lose any sleep. "Rationality is just one of the relevant factors," we say blandly; "here something else was at work."

But for better or for worse, that is how things stand. We must get used to the fact that economics is not astronomy, and game theory is not physics. We know that, in bringing up our children, we must accept each one for what *he* is, for the good that is in *him*, and must not force him into somebody else's mold. The sciences are the children of our minds; we must allow each one of them to develop naturally, and not force them into molds that are not appropriate for them.

It should be pointed out that our fields are by no means the only ones in science that are not strong on predictions and falsifiability; in which the measure of success is "does it enable me to gain insight?" rather than "what will be my observations?" Similar in this respect are disciplines like psychoanalysis, archeology, evolution, meteorology and to some extent even aerodynamics. Airplanes are not designed by solving the equations of aerodynamics: they are designed by intuition and experience, and tested in wind tunnels and in test flights. The intuition that goes into the design is based partly on theory; theory provides important general principles. But it does not do anything like predicting airworthiness of a particular design.

In the end, I think that the ordinary laws of economic activity apply to our fields as well. The world will not long support us on our say-so alone. We must be doing something right, otherwise we wouldn't find ourselves in this beautiful place today.

6 Normative Aspects of Game Theory

In the previous section we indicated in what sense game and economic theory could be considered descriptive. In this section we consider its normative aspects—game theory as engineering rather than as science.

A word of caution before starting out. The distinction between the descriptive and the normative modes is not as sharp as might appear, and often it is difficult to decide which of these two we are talking about. For example, when we use game or economic theory to analyze existing norms (e.g., law), is that descriptive or is it normative? We must also be aware that a given solution concept will often have both descriptive and normative interpretations, so that one will be talking about both aspects at the same time. Indeed, there is a sense in which the two aspects are almost tautologically the same. In the previous section we pointed out

that game theory purports to describe not *Homo sapiens*, but *Homo rationalis*; and that it actually is descriptive of *Homo sapiens* only to the extent that he can be modelled by *Homo rationalis*. On the other hand, when we come to advise people, it is clear that we should give them rational, utility-maximizing advice, i.e., precisely what *Homo rationalis* would do; so that the two aspects are in this sense quite close.

Nevertheless, if not taken too rigidly the distinction is sometimes quite useful, and we shall avail ourselves of it here.

Normative aspects of game theory may be subclassified using various dimensions. One is whether we are advising a single player (or group of players) on how to act best in order to maximize payoff to himself, if necessary at the expense of the other players; and the other is advising society as a whole (or a group of players) of reasonable ways of dividing payoff among themselves. The axis I'm talking about has the strategist (or the lawyer) at one extreme, the arbitrator (or judge) at the other. Again, the distinction is often spurious; when advising what to do, you must take into account what the other players can do, and the outcome may well be a reasonable compromise. In real life, an important function of lawyers is to restrain their clients, and to press for reasonable compromises. Conversely, reasonable compromises are based, almost by definition, on the capabilities of the players. So this distinction, too, is a blurry one; but it is useful to keep in mind.

Another distinction is between advice that is precise and/or numerical, and normative insights of a general nature. An example of the former would be a specific minimax strategy in a well defined tactical situation (e.g. a destroyer-submarine hide-and-see situation). An example of the latter would be the insight that, in negotiating an international treaty, each side should make sure that the proposed treaty represents an equilibrium point, i.e. that it is written in such a way that it is *a priori* common knowledge that it is not worthwhile for either side to violate it.

7 The Classifying Function of Game Theory

Another important function of game theory is the classification of interactive decision situations. Perhaps this could be considered part of the descriptive theory; but it really is something different, because it describes the situations themselves rather than the behavior of the participants in them. A classification of game situations is in fact as important for the normative as for the descriptive theory.

Much of science starts out with classifications; the "right" classification is often the key to a successful theory. Modern biology was made possible by Linnaeus's classification of all living things into species, genera, etc.;

geology starts with the classification of rocks; and so on. Often classifications are to some extent spurious; we know now that the notion of species is not as central as had been thought, and that life might be more accurately considered a continuum of different individuals representing different combinations of genetic characteristics. But the classification is nevertheless indispensable to enable us to organize our thoughts.

Just making the classifications is already a form of science; it enables us to relate interactive situations to each other, to identify common features, often to draw operative conclusions. We classify games into cooperative or non-cooperative according as to whether or not there is available a mechanism that can enforce agreements. The man in the street is, I think, not aware of this distinction; certainly he is not aware of its crucial importance; therefore he may think of an international treaty as if it were much like a contract between businessmen. Having classified games by their coalitional form, we become aware that the weighted majority game $[5; 2, 3, 4]$ is in a sense “the same” as $[2; 1, 1, 1]$; this is far from obvious to the man in the street.

It may seem that these applications are comparatively trivial, and that one shouldn't require the whole gigantic theoretical structure of game theory to reach this kind of elementary conclusion. Nothing could be further from the truth. It is not enough formally to define concepts like cooperative or non-cooperative in order to make them conceptually meaningful. A formal definition is given meaning and content only by experience and use. Without the theoretical work that has been done over the last decades on both cooperative and non-cooperative games, the concepts themselves would be sterile, and the fundamental distinctions they embody would be overlooked.

An interesting example of this kind of process is the distinction between games of complete and incomplete information. The early development of game theory treated only complete information games, i.e., those in which the structure of the game and the payoffs are common knowledge to the players. The restrictive nature of this assumption was not overlooked; for example, Luce and Raiffa (1957) are quite explicit on this point. Nevertheless, this was the only available theory; the result was that the existing applications all used the complete information model, for better or for worse. No adequate tool for dealing with incomplete information existed; nobody knew just how the theory would have to be modified to deal with it. After muttering a few words of apology, therefore, people simply used the complete information model as an “approximation” to the existing situation. That is, the effects of incomplete information were ignored; for all practical purposes, incomplete information games did not exist.

Then came the landmark work of Harsanyi (1967). The most important immediate effect was the realization that incomplete information games were indeed in principle no different from complete information games, that the classical game-theoretic framework applied to them as well. Without further work on the subject, it is possible that people would have continued to think mainly in terms of complete information, simply substituting a citation of Harsanyi for the muttered words of apology. But Harsanyi's work led to a stream of investigations of incomplete information models of ever-increasing volume and depth; today, incomplete information is perhaps the "hottest" topic in economic theory. All this work led to the realization that models with incomplete information are quite different from those with complete information; that they have their own problems and their own features. To cite just one example, it became clear that disagreement was a perfectly rational outcome in incomplete information cooperative bargaining. It is the sum total of the work on incomplete information models—not the mere definition—that makes the idea meaningful and gives it content. Of course, it is Harsanyi's definition that led to the development of the theory, that made the theory possible; but it is the theory itself that is decisive, not the definition.

A case in point is the ingenious experimental work of Roth and Murnighan (1982) on bargaining under complete and incomplete information, with and without assumptions of common knowledge.¹ It is interesting to compare these experiments with those of Fouraker and Siegel (1960) carried out some 20 years earlier. Fouraker and Siegel also ran bargaining experiments in both the complete and incomplete information modes. But they did not have Harsanyi's model, and so were unable to specify the incomplete information case more than that each side "is not informed" of the other side's payoffs. Roth and Murnighan, on the other hand, specified the incomplete information in terms of types, and made explicit use of the common knowledge variable (whether or not the description of the game was common knowledge).

In fact, Roth and Murnighan use little or no theory. In principle, it would have been possible to perform these experiments in 1960; in practice, this could not—and did not—happen before Harsanyi's "type" model, and the related notion of common knowledge, were not only defined, but also given a chance to mature, to become part of the atmosphere, to be *understood*. And this can happen only as a result of the development of the related theory.

1. See also Malouf and Roth (1981) and Roth, Malouf and Murnighan (1981).

8 Game Theory and Mathematical Economics as Art Forms

An alternative way of viewing game theory and mathematical economics is as art forms. The distinction between the common conceptions of science and art is in any case not sharp; perhaps our disciplines are somewhere in between. Much of art portrays the artist's subjective view of the world; art is successful when the view expressed by the artist finds an echo in the minds of his audience, when the audience empathizes with what the artist is expressing. For this to happen, the artist's statement must have some universality; it must express some insight of a general nature, must be related to the audience's experience—in brief, there must be some objectivity in it.

The case for thinking of mathematics itself as an art form is clear. Mathematics at its best possesses great beauty and harmony. The great theorems of, for example, analytic number theory are reminiscent of Baroque architecture or Baroque music, both in their intricacy and in their underlying structure and drive. Other sides of mathematics are reminiscent of modern art in their simplicity, spareness and elegance; the most lasting and important mathematical ideas are often also the simplest.

A characterization of art that I find very apt is “expression through a difficult (or resistive) medium.” (I heard this from my friend M. Brachfeld; he said he had read it somewhere, but could not remember where.) The medium may be stone or rhyme or meter or a musical instrument or canvas and paint, or the less well-defined but no less demanding medium of the novel. The resistiveness of the medium imposes a kind of discipline that enables—or perhaps forces—the artist to think carefully about what he wants to express, and then to make a clear, forthright statement.

In game theory and mathematical economics, the resistive medium is the mathematical model, with its definitions, axioms, theorems and proofs. Because we must define our terms, state our axioms and prove our theorems precisely, we are forced into a discipline of thought that is absent from, say, verbal economics.

If one thinks of mathematics as art, then one can think of pure mathematics as abstract art, like a Bach fugue or a Pollock canvas (though often even these express an emotion of some kind); whereas game theory and mathematical economics would be expressive art, like a cubist painting or Tolstoy's *War and Peace*. We strive to make statements that, while perhaps not falsifiable, do have some universality, do express some insight of a general nature; we discipline our minds through the medium of the mathematical model; and at their best, our disciplines do have beauty, simplicity, force and relevance.

9 Some Solution Concepts

The point of view I shall take in the next few sections was set forth at the end of the introductory section: that a solution concept should be judged by its performance in the applications, by the quantity and quality of the relations that it engenders, not by “armchair” philosophizing about its definition. From this point of view we will consider four of the most important solution concepts of game theory—the Nash equilibrium, the core, the N–M stable set (or “solution”) and the Shapley value. The idea is not to give an exhaustive survey of game theory or even of these four solution concepts, but to illustrate our ideas by *some* examples of *some* solution concepts. In addition to mentioning applications of these concepts, we will try to get some feel for what each of them expresses—not from its definition, but from the applications themselves.

10 The Nash Equilibrium

This is certainly the game-theoretic solution concept that is most frequently applied in economic theory. Born more than a century ago in connection with Cournot’s (1838) study of duopoly, it is now extremely common in many different applications. In perfectly competitive markets, it is closely associated with the competitive equilibrium.² Novshek, Sonnenschein and others have used it to study entry and exit.³ It is commonly used for search,⁴ location⁵ and product quality⁶ problems. In incomplete information set-ups it has been used to study auctions,⁷ insurance,⁸ principal-agent problems,⁹ again entry and exit,¹⁰ health,¹¹ higher education,¹² discrimination¹³ and a host of other particular

2. Arrow and Debreu (1954), Debreu (1952), Hurwicz (1979a, 1979b), Schmeidler (1969a, 1980), Shapley and Shubik (1967), von Neumann (1937).

3. Novshek (1980), Novshek and Sonnenschein (1978).

4. Butters (1977), Diamond (1971).

5. d’Aspremont, Gabszewicz and Thisse (1979), Hotelling (1929), Prescott and Visscher (1977).

6. Akerlof (1970), Salop (1979).

7. Griesmer and Shubik (1963), Milgrom and Weber (1982), Myerson (1981), Vickrey (1961, 1962), Wilson (1979).

8. Rothschild and Stiglitz (1976).

9. Radner (1981), Ross (1973), Shavell (1979).

10. Milgrom and Roberts (1982a, 1982b).

11. Pauly and Satterthwaite (1981).

12. Arrow (1973b), Spence (1974, 1976).

13. Arrow (1972a, 1972b, 1973a).

models. In social choice theory it is ubiquitous. It is probably safe to say that it impinges significantly on every area in which incentives are important—and this includes just about all of economic theory.

The Nash equilibrium is the embodiment of the idea that economic agents are rational; that they simultaneously act to maximize their utility. If there is any idea that can be considered *the* driving force of economic theory, that is it. Thus in a sense, Nash equilibrium embodies the most important and fundamental idea of economics, that people act in accordance with their incentives.

There is a beautiful theorem of John Harsanyi (1973) relating to mixed strategy equilibria that illuminates and underscores this idea. Mixed strategy equilibria have always been intuitively problematic because they are not “strict” (cf. the discussion of strictness below): a player will not lose if he abandons the randomization and uses instead any arbitrary one of the pure strategy components of the randomization. Harsanyi gets around this by subjecting the payoff function of each player to a slight random perturbation. Each player is privately informed of the true value of his payoff function, but the other players know only the mean. The result is that each player will be motivated to choose a particular one of the pure strategy components of his equilibrium mixed strategy. An outside observer sees the game as one of complete information, with the players playing mixed strategies. In fact, the game is of incomplete information, and the players are playing pure strategies. The mixed strategy models the ignorance of the outside observer and of the other players, not a conscious randomization.

The beauty of this model is that it rings so true. No game can really be of complete information; there are always nuances in tastes and whims of the other players of which we cannot be aware.

Nash equilibrium has spawned a number of variants and related concepts. In the applications the most significant of these to date have been Selten’s “perfection” concepts: first the subgame variety (Selten, 1965), later the more general “trembling hand” variety (Selten, 1975). Many of the economic applications mentioned above in fact use perfect rather than ordinary equilibrium points (EPs). Related to perfection are the notions of subgame symmetry proposed by Kalai and Samet (1985), Selten (1980) and others, the Kreps–Wilson (1982a) sequential equilibrium, Kalai and Samet’s (1984) persistency, Myerson’s (1978) properness, and recent concepts developed by Kohlberg and Mertens (1986) which are based on topological properties of the equilibrium correspondence.

The *definitions* of these concepts usually express a certain robustness or continuity; if the game is changed a little, the equilibrium won’t change

much. In *practice*, they express a more subtle idea, a sort of forward-lookingness. Nash equilibrium may take threats into account; i.e., it allows the players to knuckle under to threats even when it would hurt the threatener to carry them out. The above concepts do not; in a sense, they say “no matter what happens, do what is best for the future, and assume that it is common knowledge that everybody will do so.” The problem remains as to how to resolve uncertainties about the past that affect optimal behavior in the future; this is where the above concepts differ, the “forward-lookingness” being a kind of common denominator.

Several other kinds of variants of the Nash concept appear in the literature. An equilibrium is *strict* if the inequalities appearing in its definition are strict, i.e., if each player’s equilibrium strategy is the *unique* best reply to the other players’ strategies. This implies stability of the equilibrium; i.e., if some or all players use strategies slightly different from the equilibrium strategies, all are motivated to return to the equilibrium. It thus expresses robustness in a different sense than that expressed by perfection: robustness as a function of the strategies rather than of the game. Harsanyi (1974, 1975) and Selten have used this kind of stability, together with Pareto domination and many other ideas, in constructing an elaborate theory that chooses a unique equilibrium point for each game. This theory has been applied to several specific game-theoretic models, most of them bargaining models with either complete or incomplete information (e.g., Selten and Güth, 1982).

An equilibrium is *strong* if no coalition of players can all gain by a simultaneous deviation (while the players outside the coalition maintain their strategies). Strong equilibria have significant applications in repeated games, as we will see below; they are also important in social choice theory, where they appear, e.g., in the context of manipulations of elections.

Two other variants that have been proposed are the subjective and the correlated equilibrium (Aumann, 1974). In subjective equilibria, the players use uncertain events for which they have differing subjective probabilities to do their randomizing; in correlated equilibria, the randomizations of the players need not be independent. Correlated equilibria have been applied to various contexts, including repeated games (Forges, 1985, 1986). Subjective equilibria have to date found little direct application to specific economic or other models, though various general theorems have been proved relating them to objective equilibria. But there has been one very interesting spinoff: they led to the notion of *common knowledge*, and to the subsequent development of the theory in this area. This notion was originally defined by the philosopher Lewis (1969) and

found its way into the decision sciences as a result of an investigation of the properties of subjective equilibria (Aumann, 1976).

Let us now return to the applications. A surprising and beautiful application of Nash equilibrium points is to evolutionary biology (cf. the discussion in section 3 above). Maynard Smith's (1982) *Evolutionarily Stable Strategy* (ESS) represents a kind of Nash equilibrium point. Recently Selten (1980) has combined Maynard Smith's ESS with the perfection notion and developed a comprehensive game-theoretic approach to animal conflicts and similar biological contexts. This has not remained in the purely theoretical sphere; there are applications to empirical investigations of the behavior of specific animals (speckled wood butterflies, etc.).

Another wide area of application that we have not yet mentioned is that of cooperative games. It will be recalled that cooperative games differ from non-cooperative games only in that, in the former, agreements can be enforced. In his 1951 *Annals of Mathematics* paper, Nash suggested that EPs could be applied to cooperative games if the pre-play bargaining procedure by which players form coalitions and agree on payoffs were formalized, and made a proper part of the game. Once this is done it is no longer necessary to specify that the players may make enforceable agreements, since these are possible within the framework of the game. The game is therefore non-cooperative, and may be analyzed by means of equilibrium points.

For a time this program had only limited success. In most bargaining procedures, a very large number of equilibrium points is possible; there may also be considerable dependence on the specific form of the bargaining procedure. Nash himself (1950) proposed a way of getting around some of these problems in the case of his two-person bargaining problem, but it was fraught with difficulties and he never managed to carry it out explicitly.

As often happens, the first successes in this area were indirect. Rather than using specific bargaining procedures, repeated games were used as a paradigm for bargaining procedures. Underlying and justifying this approach is the famous "Folk Theorem," according to which *any* individually rational feasible outcome of a game can be achieved as an equilibrium outcome of an infinite repetition of that game; as shown by Rubinstein (1995) and others, this holds also for perfect equilibria. As early as 1959 it was shown that strong equilibria of the repeated games are associated with the core of the one-shot game (Aumann, 1959).

Later attention shifted to incomplete information repeated games. In such games, bargaining involves a subtle interplay of concealing and

revealing information: concealing, to prevent the other players from using the information to your disadvantage; revealing, to use the information yourself, and to permit the other players to use it to your advantage. In order to study these phenomena, attention was first concentrated on two-person zero-sum repeated games. A large, subtle and deep literature on this subject was spawned, which spilled over into related fields such as stochastic games. Two-person zero-sum games certainly cannot give us any insight into bargaining, but here again we see the importance of not taking too narrow an approach. In order to understand cooperative games of incomplete information, we study non-zero-sum repeated games of incomplete information; and for this purpose, it is necessary first to understand zero-sum repeated games of incomplete information. After zero-sum repeated games were studied intensively for about 15 years, a major breakthrough in the area of non-zero-sum repeated games was recently achieved by Sergiu Hart (1985), who gave a complete characterization of equilibrium outcomes in such games when there are just two players and there is complete information on one side. The characterization is not just technical; it gives important qualitative insights into the nature of incomplete information bargaining.

Complete information repeated games have also been taken up again in recent years. Important applications of equilibrium points in repeated games have been made in economic contexts such as altruism, principal-agent problems, insurance, oligopoly and the development of reputations by several authors.¹⁴ The repeated prisoner's dilemma has always been an object of considerable theoretical as well as experimental interest; this has led in recent years to the development of several additional variants on the equilibrium concept. Worthy of special mention here is Axelrod's (1984) beautiful computer experiment.

Quite recently it has been shown that the more direct approach to the Nash program for attacking cooperative games can prove fruitful after all. Using a bargaining procedure in which time is costly, Rubinstein (1982) showed that perfect equilibrium outcomes lead to equal division in a two-person bargaining problem with a flat efficient frontier; later Binmore (1987) extended Rubinstein's result by showing that, in any two-person bargaining problem, Nash equilibrium outcomes are associated with Nash's product maximization solution (which coincides with the non-transferable utility value). The Rubinstein-Binmore results are particularly beautiful because they express an important practical insight:

14. Friedman (1977), Hammond (1975), Kreps, Milgrom, Roberts and Wilson (1982), Kreps and Wilson (1982b), Kurz (1978), Milgrom and Roberts (1982b), Radner (1980, 1981), Rosenthal (1979), Rubinstein (1979) and others.

impatience is an important component in reaching a compromise; the players are more likely to reach a reasonable compromise if delay is costly to them, if the value of the product that is being bargained for deteriorates with time.

11 Nash Equilibrium: Summary and Conclusions

Nash equilibrium is without a doubt the most “successful”—i.e., widely used and applied—solution concept of game theory. It touches almost every area of economic theory, as well as social choice, politics and many other areas of application. Within game theory itself it engenders a host of relationships. Basically a non-cooperative concept, it has nevertheless been applied with considerable success to cooperative models.

Conceptually, the Nash equilibrium and most of its variants express the idea that each player individually maximizes his utility; it is a simple expression of the rationality of the individual player. Two of its variants, though—the strong equilibrium and the Harsanyi–Selten unique equilibrium—go beyond this to express some form of cooperation or joint rationality.

The definition of the Nash equilibrium is in form extremely simple. Moreover, the concept is mathematically very tractable and easy to work with. But there are problems with its intuitive interpretation. In games of perfect information, perfect equilibria can be arrived at by a sort of dynamic-programming, backwards-induction procedure whose intuitive content is very clear and compelling. In other games it is by no means clear how the players would arrive at an equilibrium, why they should play equilibrium strategies, and how a specific equilibrium would be chosen from among the set of all equilibria. There are indeed games in which the Nash equilibrium looks very strange and counterintuitive. For a long time it was thought that a Nash equilibrium could be thought of as a self-enforcing agreement, but recently it has been shown that this, too, is incorrect: there are games with multiple Nash equilibria, in which an *agreement* to play a certain one of them does not increase the chance of its *actually* being played (Aumann, 1990).

Nash equilibrium therefore is an example *par excellence* of our basic thesis. On the one hand, philosophical analysis of the definition itself leads to difficulties, and it has its share of counterintuitive examples. On the other hand, it is conceptually simple and attractive, and mathematically easy to work with. As a result, it has led to many important insights in the applications, and has illuminated and established relations between many different aspects of interactive decision situations. It is these applications and insights that lend it validity.

12 The Core

An outcome x of a game *dominates* another outcome y if there is a coalition S that can achieve x by its own efforts, each of whose members prefers x to y . The *core* of the game is the set of all undominated outcomes.

The “core” applies to cooperative games; it is the cooperative solution concept that is perhaps best known to economists. Most famous among its applications is the core equivalence principle, which states that the core coincides with the set of competitive (price equilibrium or Walras) outcomes in perfectly competitive markets with many traders, each individual one of whom is insignificant. First demonstrated by Edgeworth (1881) more than a century ago, it was forgotten, then exhumed and reformulated by Shubik (1959) about 25 years ago. Since then an amazingly rich and deep literature has sprung up, all of it focusing on this one basic principle. It has been expressed as a limit theorem¹⁵ for markets with n traders as $n \rightarrow \infty$, in the context of a non-atomic continuum of traders,¹⁶ and in the context of non-standard analysis, in which individual traders formally appear as infinitesimal.¹⁷ Among the aspects that have been discussed in one or more of these forms or modes are rates of convergence,¹⁸ extension to productive economies,¹⁹ continuity as the utility functions or initial bundles vary,²⁰ and the extent to which the result remains true when one restricts the family of “permitted” coalitions (coalitions via which domination may take place).²¹ Considerable effort has also gone into carefully determining the boundaries of the core equivalence principle, i.e., just where it ceases to hold. For example, in general, the principle no longer holds when there is a continuum of goods (qualities, locations) as well as traders.²²

It should be mentioned that in any market—even one with a small number of traders—each competitive outcome is in the core. It is the converse that requires a large number of individually insignificant traders.

While the lion’s share of the literature on the core has been devoted to the equivalence principle, there has been significant work on other appli-

15. Debreu and Scarf (1963).

16. Aumann (1964), Vind (1964).

17. Brown and Robinson (1975).

18. Anderson (1978), Arrow and Hahn (1971), Aumann (1979), Bewley (1973), Cheng (1981), Debreu (1975), Grodal (1975), Shapley (1975).

19. Hildenbrand (1974), Sondermann (1974).

20. Dierker and Dierker (1972), Kannai (1970, 1972).

21. Grodal (1972), Mas-Colell (1977), Schmeidler (1972), Vind (1972).

22. See Ostroy (1981).

cations. One area in which much work has been done is oligopoly or syndication; in the continuum framework, one would express this by saying that the space of traders has atoms (the oligopolists or syndicates) as well as a non-atomic part (the “small” agents).²³ One typical result is that, when all the oligopolists are similar in utility and endowment (although they may differ in size), the equivalence principle continues to hold (Shitovitz, 1973). This relates to Bertrand’s (1883) classical approach to duopoly theory, in which the duopolists compete until a perfectly competitive outcome is reached.

The core has also been applied to public goods; it has been shown that the Lindahl equilibrium is always in the core (Foley, 1970). The equivalence principle, however, does not hold in public goods economies with a continuum of agents; there may be core points that are not Lindahl equilibria.

An area that is currently very active is the study of the core in discrete markets (e.g. private homes) with a fixed finite number of traders.²⁴ The classic paper in this area is Gale and Shapley’s (1962) “College Admissions and the Stability of Marriage.” The main issue in those cases is the non-emptiness of the core; there need be no price equilibrium in these markets, and the core becomes the expression of competition.

There is no general existence (or rather non-emptiness) theorem for the core, and if one strays too far from classical market models, it is indeed often empty. In voting games, for example, the core is always empty unless there are veto players (in which case the core shares all the payoff among the veto players). Economies with S-shaped production curves (initially increasing returns, then decreasing) also have empty cores. Non-emptiness of the core expresses a situation in which there is no disincentive for the all-player coalition to form, in which each set of players can do at least as well in the framework of the all-player coalition as it could by itself—a kind of consumer surplus from the formation of the all-player coalition. In transferable utility games v , it is related to superadditivity of the coalitional worth function, or more specifically to local superadditivity “at” the all-player coalition N (i.e. the worth of N is at least as great as the sum of the worths of the coalitions in any partition of N). In fact, the Bondareva–Shapley necessary and sufficient condition for the non-emptiness of the core (“balancedness”) can be thought of as a

23. Aumann (1973), Dreze, Gepts and Gabszewicz (1969), Gabszewicz and Dreze (1971), Gabszewicz and Mertens (1971), Shitovitz (1974).

24. Crawford and Knoer (1981), Dubins and Freedman (1981), Gale (1984), Kaneko (1976), Knuth (1976), Quinzii (1984), Roth (1982, 1984), Shapley and Scarf (1974), Shapley and Shubik (1972), Thompson (1980).

strong kind of superadditivity at N , as follows. Suppose we extend the game to include “part-time” coalitions. That is, each player may split his time into several parts, each of which he devotes to a different coalition; if the players in a coalition S each devote a proportion θ of their time to S , then the resulting “part-time” coalition θS has worth $\theta v(S)$. Then the core is non-empty if and only if the worth of N is at least as great as the sum of the worths in any partition of N into “part-time” coalitions (Bondareva, 1962, 1963; Shapley, 1967).

We noted above the loose association between non-empty cores and markets. In the case of transferable utility (TU) games this was made precise by Shapley and Shubik (1969): a game is associated with a market if and only if it *and* all its subgames have non-empty cores. And this happens if and only if the *extended* game is super-additive in the ordinary sense; i.e., the worth of *each* (part- or full-time) coalition S is at least as great as the sum of the worths of the (part- or full-time) coalitions in any partition of S . (This is what is called a “totally balanced game.”)

As far as non-transferable utility (NTU) games are concerned, the situation is roughly as follows. H. Scarf (1967) has extended the notion of balancedness to such games, has shown that balanced games have non-empty cores (but not conversely), and that markets are balanced. Billera and Bixby (1973a, 1973b, 1974), Mas-Colell (1975) and others have extended the characterization of market games to the NTU case; the subject is, however, a difficult one, and the results are not as complete as in the TU case.

One cannot complete the discussion of the core without mentioning computation. Using his definition of balancedness, Scarf (1967) developed an algorithm for finding a point in the core. This algorithm was the forerunner of the algorithms later developed by Scarf (1973) for finding competitive equilibria and, more generally, fixed points of mappings. Here again is an unexpected spinoff: investigation of the core eventually led to the creation of a whole new branch of numerical mathematics, the calculation of fixed points.

Like the Nash equilibrium, the core is subject to puzzling counter-intuitive examples. In a market with two complementary goods—e.g. right and left gloves—suppose m traders are endowed with one right glove each and $m + 1$ traders with one left glove each: then the core consists of a unique point, under which the owners of the left gloves must simply give all their merchandise, for nothing, to the owners of the right gloves. This might perhaps be viewed as an extreme expression of cut-throat competition.

More puzzling, though, is the case in which two traders hold one right glove each, one trader holds a left glove, and one trader holds *two* left

gloves. In this case, too, the holders of the left gloves must give their merchandise, for nothing, to the owners of the right gloves. This is more difficult to understand: by the simple expedient of throwing away a glove—an action that he can take by himself, without consulting anybody—the holder of the two left gloves can make the situation completely symmetric. Yet the unique core point assigns him nothing.

We have already mentioned that the core (more precisely, the β -core) of a game coincides with the strong equilibrium payoffs in an infinite repetition of that game. Some relations with other solution concepts are as follows. The core is included in each von Neumann–Morgenstern (N–M) stable set (see section 14) as well as in the bargaining set (Davis and Maschler, 1967; Peleg, 1967); when non-empty it intersects the kernel (Davis and Maschler, 1965) and contains the nucleolus (Schmeidler, 1969b). There are in fact some beautiful geometric characterizations of the kernel and the nucleolus as “central” points in the core (Maschler, Peleg and Shapley, 1972). There is no clear general relationship with the value; but in non-atomic market games the core contains the value, and coincides with it if the market is sufficiently smooth (Aumann, 1975). A convex TU game always has a non-empty core (Shapley, 1971); it contains the value, includes the kernel (which coincides with the nucleolus for such games) and coincides with the bargaining set and the N–M stable set (there is only one stable set in such games) (Maschler, Peleg and Shapley, 1972). From all this there emerges a picture of the core as being a “central” solution concept, one with significant relations with many other concepts. Conceptually, too, one may think of the core as a sort of starting point or first approximation, with other solution concepts overcoming shortcomings in the core in various different ways.

13 The Core: Summary and Conclusions

Most applications of the core are to economic contexts, more specifically to market contexts of one form or another. The outstanding application is embodied in the core equivalence principle, which relates the core of a perfectly competitive market to its competitive equilibria.

Conceptually, the core expresses the idea of unbridled competition; non-emptiness of the core expresses the idea that such competition can lead to stability, that there is an outcome consistent with it. In practice, this happens chiefly in economic contexts of the kind described above. Political contexts are inherently less stable, and for them the core is often empty; they are covered by concepts such as the value, which expresses compromise or average outcome, or the von Neumann–Morgenstern (N–M) stable set, which expresses a weaker and more complex kind of

stability than the core. In a game with both political and economic aspects, the core often is not sensitive to the political aspects even when it is not empty (e.g., in the voting model of Aumann and Kurz, 1977, or, more generally, Neyman's (1985) political-economic games).

It should be stressed that this image of the core as expressing competition emerges from the applications; it is not in any sense obvious from the definition. The relationship between the core and competition is indeed a two-way street; several investigations of the core—most notably the equivalence principle itself, but also Ostroy's investigations of the equivalence principle with a continuum of goods, or the study of the core in one-homogeneous non-atomic superadditive games—have shed light on the nature of competition, and have sharpened our idea of what we mean when we talk about a perfectly competitive set-up.

The definition of the core is in form extremely simple, and the concept is mathematically fairly tractable and easy to work with. Intuitively, it is perhaps the clearest and most transparent concept in the theory of games. Its main fault is that it is often empty; therefore, unlike the Nash equilibrium, the von Neumann–Morgenstern stable set and the value, it cannot serve as a unifying theory for the rational social sciences, or even for all of economics. It is closely associated with the competitive side of economics, and there it gives important insights.

There are important relationships between the core and almost all other solution concepts. In a sense, the extremely simple and transparent definition of the core may be considered a starting point for the reasoning leading to more sophisticated concepts like the N–M stable set, the nucleolus, and so on. Thus, in more senses than one, the core occupies a central position in cooperative game theory.

14 The Von Neumann–Morgenstern Stable Set

The definition of the core of a game as the set of all undominated outcomes is subject to the following conceptual query. Suppose we think of outcomes in the core as “good” or “stable.” Then we should not exclude an outcome just because it is dominated by *some* other outcome; we should demand that the dominating outcome itself be “stable.” If the outcome y is dominated by an outcome x that is not itself “stable,” then the argument for excluding y is rather weak; proponents of y can argue with some justice that replacing it with x would not lead to a more stable situation, so we may as well stay where we are. If the core were the set of all outcomes not dominated by any element of the core, then there would be no difficulty: this, however, is not the case.

We are led to the following definition: a set K of outcomes is called *stable* if it is the set of all outcomes not dominated by any element of K . The stable sets are precisely the *solutions* of von Neumann and Morgenstern (1944).

In two-person bargaining games, there is a unique stable set, which consists of all efficient (i.e., Pareto-optimal) and individually rational outcomes. Usually, however, there is more than one stable set. For example, in the three-person majority game there is a unique symmetric stable set; it consists of three outcomes, each of which provides for formation of one of the three possible two-person coalitions, which then divides the available payoff equally. But there are many non-symmetric stable sets, called “discriminatory.” Each of these specifies that some fixed two-person coalition forms, assigns to the remaining player a fixed amount (which must be less than half the total payoff), then divides the remaining amount in an arbitrary fashion between the two players in the coalition. Geometrically, therefore, each such stable set is an interval.

Intuitively, the symmetric solution represents a situation in which, *a priori*, each of the three players is a possible coalition partner; which coalition will form is determined by negotiations between the players, based on payoff. The three outcomes in the solution are stable as a set, not individually; in a sense, each is “right” only because the others are also there. Von Neumann and Morgenstern use the term “standard of behavior”; the idea is that, since each coalition knows or expects that the other coalitions will divide 50–50 if they form, it too is motivated to divide 50–50. It represents a form of social organization that, together with corresponding norms of dividing the payoff, is stable as a whole.

The discriminatory solutions also represent stable forms of social organization, of a different kind. One player is excluded *a priori* from the negotiations. He may or may not be given a certain sum “to keep him quiet,” but this sum is fixed and not subject to negotiation, and neither of the other players considers it a real possibility that the excluded player will enter the negotiations. As a consequence, there is no constraint on them; the negotiations between them turn into a two-person bargaining game, and the solution of this, as we have seen, is the entire efficient individually rational interval.

Consider next the “glove market” of section 12. Take the very simple case in which $m = 1$; i.e., there are just three players, two holding left gloves and one holding a right glove. It is convenient to think of the former as “sellers,” the latter as a “buyer” (of left gloves). The core says that the sellers simply give their merchandise to the buyer: they get nothing out of it. This corresponds to unbridled competition, in which each seller tries to underbid the other. The N–M stable sets, on the other

hand, have the two sellers colluding, i.e., forming a single bargaining unit, to bargain with the buyer. Between this bargaining unit and the buyer we again have a two-person bargaining game, which leads to an interval. How the two sellers divide what they get depends on the particular stable set under consideration. Some stable sets have them dividing it in a fixed proportion, e.g., 50–50; in others they may divide 50–50 up to a certain amount, with the remainder going to the first seller; and so on. There are many possible arrangements. The important thing is that the sellers must agree beforehand on a definite way of dividing whatever they get, and it must be worthwhile for each seller that the coalition of sellers get as much as possible from the buyer, so that there is no possibility of the sellers entering into competition with each other during negotiations with the buyer.

Note that, unlike in the three-person majority game, in this game all the stable sets involve essentially the same form of social organization: the sellers organize in a unit that bargains with the buyer. The different stable sets differ only in the internal arrangement that the sellers may have for dividing what they get from the buyer.

A good chunk of stable set theory is devoted to weighted majority voting games. These may be thought of as parliaments; the “players” are the parties rather than individual members. The interesting cases are those in which no single party has a majority. Such a game is called *homogeneous* if each minimal winning coalition has total “weight” equal to the “quota” (the “weight” of a party is the number of members it has; the “quota” is the minimum total weight needed to “win”—i.e., to form a government or pass a bill). It is called *strong* if the quota is a simple majority and the total weight is odd. It would seem that strong homogeneous voting games are few and far between; but in fact there are many games that are equivalent to strong homogeneous games, though this is not apparent from the given numbers. For example, this is the case for every strong game with fewer than six players (von Neumann and Morgenstern, 1944).

Every strong homogeneous game has a stable set that predicts that some minimal winning coalition will form, and will divide the payoff among the parties in the coalition in proportion to their representation in parliament. This generalizes the symmetric solution in the three-person majority game.

When the game is not strong, as for example in American political conventions in the first half of the century, the social organization predicated by stable set theory is not in terms of minimal winning coalitions, but in terms of minimal blocking coalitions. A coalition S is *blocking* if

it can prevent a win, i.e., if each winning coalition must intersect S . In the convention example, any coalition with more than one-third of the members is blocking.

This is a nice example of the kind of insight that game theory affords. All the evidence from strong games—and this includes far more than the simple homogeneous weighted majority case cited above—indicated that the important coalitions were the winning ones. Then came Bott's (1953) analysis of the simplest symmetric voting games that are not strong, which pointed to the blocking coalitions, and indicated that in strong games the winning coalitions are important because they coincide with the blocking ones. As soon as one realizes that in non-strong games the players organize themselves into blocking coalitions, one understands much better what went on in those political conventions, the deals that were made in the smoke-filled back rooms, the payoffs to special interest groups that were, in fact, blocking coalitions.

In brief, the outcome is unexpected, but once there, it sounds right and natural; it clicks into place. This is game theory at its best.

The literature is full of dozens of applications of the N–M theory, each with its own stable set or class of stable sets, each with its own story and interpretation involving some particular form of social organization. There is no space here even to mention all these results; indeed, we are not trying to survey the field, but only to give the flavor of the kind of insight that is afforded by the theory.

We mention only one more outstanding application, of a more recent vintage than the foregoing ones. This is S. Hart's (1974) analysis of the formation of cartels in large markets. The market has a continuum of traders, divided into disjoint types (traders of the same type have the same endowments and utility functions). Each type has a monopoly on some commodity; more precisely, each commodity is held by only one type. Then the symmetric stable sets of the market correspond to the stable sets of the market in which there is only one trader of each type. (A stable set is *symmetric* if it treats traders of the same type in the same way, i.e., if it is invariant under permutations of each type separately. It is not required *a priori* that individual outcomes be symmetric.)

The interpretation, of course, is that the types organize themselves into cartels, each of which acts as a single player.

Consider, for example, a *two-type* market of this kind. Hart's theorem says that each of the two types will bargain as a unit. The result is a two-person bargaining game, which as we know has only one solution, consisting of all efficient individually rational outcomes. In the original game, therefore, there is a unique symmetric solution, consisting of all

efficient individually rational outcomes that assign the same bundle to traders of the same type. In contrast, the core of the two-type non-atomic market is very different from that of its two-person cousin; in the non-atomic case it contains only the competitive outcome (or outcomes), in the two-person case it contains the whole bargaining range.

Most of stable set theory deals with TU games; Hart's theorem is one of the few instances of a successful NTU analysis.

There is no general existence theorem for stable sets. The existence problem was open for many years. R. Stearns (1965) found a seven-person NTU game without a stable set, and W. Lucas (1969) made headlines by finding a ten-person TU game without a stable set. Recently Lucas and Rabie (1982) have found a coreless 14-person TU game without a stable set.

But non-existence is not nearly as serious a problem for stable sets as for the core. The counterexamples are ingenious, difficult and deep; but there is no question that they are contrived. They do not appear to correspond to any economic, political or social reality; there is no broad class of games known to have no stable sets. Lucas's examples are of great importance because they show that it is hopeless to look for a general existence theorem; in practice, though, one can usually find a stable set if one tries hard enough.

15 N–M Stable Sets: Summary and Conclusion

Conceptually, stable sets express the idea of social organization—minimal winning coalitions, blocking coalitions, systematic discrimination against players or groups of players, cartel formation, groups within groups and so on. The organizational forms are often of great subtlety; they are endogenous, they spring from the analysis. This is quite different from other game-theoretic approaches to social organization, in which the forms (e.g., partitions into coalitions) are exogenous, and one looks for a structure of the given form possessing one or another kind of stability. In this sense the N–M theory is of a subtlety and depth unparalleled in game theory.

As in the case of the core, it should be stressed that this image of stable sets emerges from the applications. It is not in any sense obvious from the definition; indeed, there is no hint of organizational forms in the definition—the definition refers only to outcomes. The conceptual connection with organizational forms is in the interpretation only and has no formal status; and it is a matter of some mystery that this particular definition should lead so frequently and consistently to this kind of interpretation.

A stable set is stable as a whole; unlike core outcomes or Nash equilibria, individual outcomes in a stable set possess little or no stability by themselves. This, too, emerges from the applications; one would not have expected it from the definition as such, which in form very much resembles that of the core.

Stable set theory has been moderately “successful.” It has a large number of applications. They are not limited to one area, but are spread over the whole gamut of applications of game theory—political, economic, production with constant or increasing or S-shaped returns (Hart, 1973) and so on. As we have seen, the applications often yield important insights; and while the insights usually have to do with organizational forms, they are, other than that, often quite different from each other. In an important sense, therefore, the N–M theory can be considered unifying.

What has prevented the theory from being more successful is that it is so difficult to work with. Finding stable sets involves a new *tour de force* of mathematical reasoning for each game or class of games that is considered. Other than a small number of very elementary truisms (e.g., that the core is contained in every stable set), there is no theory, no tools, certainly no algorithm. There are perhaps some rough methods, ideas that recur here and there, tricks of the trade; people with experience are better at it than newcomers. But basically you just have to slug it out anew every time. And because stable sets do not always exist, you cannot even be sure that you are looking for something that is there. NTU games are even more difficult to solve than TU games. This has led to the paucity of NTU results; and, because of the central importance of NTU games in the applications, this is another factor limiting the “success” of the theory.

But while there is no mathematical theory that will help one to find stable sets, there is a sort of qualitative classificatory theory in the interpretations. Similar qualitative phenomena pop up in solutions of many different games. Minimal winning and minimal blocking coalitions occur in many stable sets; in interpreting the solutions of a three-person market game we make use of the solution of two-person bargaining games; “bargaining curves” appear in stable sets; and so on. Other than hinting at what to look for, this perhaps is not of much use in finding stable sets. But it does lend coherence and unity to the qualitative theory.

16 The Shapley Value

Like the core and the N–M stable set, the Shapley value (Shapley, 1953) is a cooperative solution concept; i.e., it applies to cooperative games

only. In finite TU games, it assigns a unique outcome to each game, which can be thought of as a sort of average or expected outcome, or an *a priori* measure of power. Other interpretations will be discussed below.

The Shapley value has a very broad spectrum of applications. Voting games have been studied thoroughly. In the UN Security Council, the “Big Five” hold more than 98 per cent of the total power (i.e., value) (Shapley and Shubik, 1954). In a parliament with one large party and many small parties, the large party holds a disproportionately large proportion of the power (Shapiro and Shapley, 1978; Milnor and Shapley, 1978). For example, if the large party has more than half the seats, it has all the power; if it has one-third of the seats, it has approximately half the power. In the American Electoral College, the large states again have a disproportionately large share of the power; a state like California, for example, has about 10 per cent more power than could be expected from its proportion in the Electoral College (Riker and Shapley, 1968).

These phenomena correspond well with our experience, and also with our intuitions (which are, of course, the creation of our experience). Perhaps the disproportionate strength of the Big Five in the UN Security Council comes as somewhat of a surprise to the casual observer; but the careful student will have realized that nothing can be done without them, and once they do all agree, then enough of the other countries will also. That a party with a majority in parliament has all the power is obvious. Perhaps the statement about a party with a third of the seats having half the power is less obvious; but it will not sound so strange to some of the older Israelis, who may remember the disproportionate power of the Labor Party in the long years when it held only one-third of the seats in parliament. Altogether this corresponds nicely to the dictum “Unity is power” (“L’Union fait la force”). As for the Electoral College, presidential candidates spend most of their campaign time in “key states,” and these rarely include Nevada.

Let us next consider a parliament with *two* large parties and many small ones. Suppose, for example, that each of the large parties has one-third of the seats, the remaining seats being spread among many small parties. Then unity is *weakness*! The large parties have only about one-quarter of the power each, and the small parties have half the power between them (Milnor and Shapley, 1978). It would be foolish for the small parties to form a “common front.”

This may sound paradoxical, but, again, students of the Israeli scene will not be surprised. In the presence of the two large parties—the Likud and Labor—there seems to be no tendency for the smaller parties (e.g., the four religious parties) to form a common front.

This curious phenomenon can be understood as follows. If the two large parties form a coalition with each other, they will have between them two-thirds of the vote, which is much more than needed, and they will have to share power equally. There is therefore a tendency for the large parties to seek coalition partners among the smaller parties; only half the smaller parties will be needed. However, competition will develop among the two large parties for the favors of the small parties; this will drive up the “price” of the latter until the point is reached where it is hardly more worthwhile for either large party to go with small parties than for it to join forces with the other large party (as has, in fact, happened in Israel).

The “competitive price of a vote” is an idea that appears in the intuitive interpretation of results about the value in several voting contexts. Another example is in voting for non-exclusive public goods, in which under certain circumstances the choice of public goods is not affected by who is allowed to vote; what is happening is that, because of the non-exclusivity, the “price of a vote” is next to nothing, and so the vote becomes insignificant (Aumann, Kurz and Neyman, 1983, 1987).

Normatively, values have been used to analyze weighted voting in connection with the US Supreme Court “one-man, one-vote” decision. According to this decision, equal representation in the legislature for districts with substantially different populations is unconstitutional. Rather than redistricting frequently, it was proposed to use weighted voting. Theoretically, this could lead to a situation in which a single district has a majority in the legislature; and even when this does not happen, the power conferred by weighted voting on the various representatives is by no means proportional to their weight, as we have seen. An analysis of this situation in terms of various philosophies of legislative government was carried out by Riker and Shapley (1968): it turns out that it makes a considerable difference whether the “players” are the districts, or the individual voters in the district. Testimony on values has been an important factor in deciding court cases on re-districting.²⁵

Let us turn to economics. In large markets, the value consists of competitive outcomes. In the smooth case (sufficiently differentiable utilities, no corner solutions), it consists of *all* competitive outcomes; i.e., it coincides with the core. In non-smooth situations the core may be quite large; in these cases, the value occupies a central position in the core. For example, in large TU markets in which the core has a center of symme-

25. Testimony on values was heard by the Supreme Court during the following cases: Baker v. Carr, 369 US 186 [1962], Gray v. Sanders, 372 US 386 [1963], Wesberry v. Sanders, 376 US 1 [1964], Reynolds v. Sims, 377 US 533 [1964].

try, the value is that center of symmetry (Hart, 1977a, 1977b). When the core has no center of symmetry, the value is obtained by calculating, for each trader, the core point that is best for him, then averaging over all traders (Hart, 1980).

Values have also been used to analyze problems of taxation,²⁶ public goods,²⁷ monopoly,²⁸ increasing returns²⁹ and other specific economic models.

In two-person NTU games, the value coincides with the bargaining solution of Nash (1950). If one uses Harsanyi's (1959) method for deriving the coalitional worth ("characteristic function") from the strategic ("normal") form, this leads to Nash's solution (1953) for two-person cooperative games.

It is worthwhile here to point to an interesting normative implication of Nash's bargaining model, as well as of other bargaining models. A fundamental feature of these models is that, while threats are never carried out (in the complete information case), the ability to make them plays a crucial role in determining the final Pareto-optimal accommodation between the bargainers. Our ability to fight a war determines how we shall live at peace. It is a mistake to think that atomic weapons, for example, exist only for defense in case of war: on the contrary, they determine how we live on a day-to-day basis.

A surprising normative application of value theory, which has recently gathered quite a bit of momentum, is to cost allocation. Specific applications include, for example, water management,³⁰ electric power,³¹ pollution treatment,³² intra-company transfer prices,³³ allocation of TVA projects,³⁴ allocation of taxes,³⁵ public utility pricing³⁶ and water resource development.³⁷ Another application, spectacular because of its depth and complexity, is to internal pricing of long-distance telephone

26. Aumann and Kurz (1977), Gardner (1984), Imai (1983), Kurz (1977, 1980), Osborne (1984).

27. Aumann, Kurz and Neyman (1983, 1987), Gardner (1975), Harsanyi (1959), Rosenmuller (1981, 1982).

28. Guesnerie (1977).

29. Mas-Colell (1980).

30. Bogardi and Szidarovsky (1978).

31. Gately (1974).

32. Loehman, Pengry and Whinston (1973).

33. Manes and Verrecchia (1982), Shubik (1963).

34. Parker (1943), Ransmeier (1942), Straffin and Heaney (1981).

35. Scott (1981), Verrecchia (1981).

36. Sharkey (1982).

37. Suzuki and Nakayama (1976).

calls in a large organization; a system of this kind, developed by Billera, Heath and Raanan (1978), was adopted by Cornell University about five years ago.

We illustrate the method with the airport landing application, due to Littlechild and Owen (1973). The runway at an airport must be large enough to accommodate the largest aircraft that will land there. But it would be absurd to charge the same fees to a Boeing 747 and a Piper. The system used at many airports is as follows. One figures the cost of a runway sufficient for the smallest plane that lands at the airport; this cost is divided equally between all landings. Next, one figures the cost *increment* necessary to build a runway long enough for the next largest plane; this increment is divided equally between all landings *except* those of the smallest planes. And so on.

Littlechild noticed that this is equivalent to charging each landing the Shapley value of a certain game. The players are individual landings; the worth of a coalition is the cost of building an airport that would accommodate those landings. The extension to more complex cost allocation problems is clear, at least in principle.

In the 2000-year-old Babylonian Talmud, the Shapley value crops up in the solution of a bankruptcy problem. Of course, the Talmud does not justify this solution by means of axioms; it does not have the general notion of a game as we know it. The problem the Talmud deals with is identical in mathematical form to the landing fee problem; its solution has independent intuitive appeal, and presumably the Talmud's choice of solution is based on this. (This is confirmed by medieval commentators.) The interesting point is that the Shapley value is able to unify so many different, apparently disparate phenomena.

Incidentally, it is this kind of legal application that I referred to in section 6 as being somewhere between descriptive and normative. It is *descriptive* of existing legal *norms*.

The Shapley value was originally axiomatized for finite TU games (Shapley, 1953); the axiomatization was later extended to games with a continuum of players (Aumann and Shapley, 1974) and to NTU games. A. Roth (1977) used the original TU axiomatization to give a rigorous treatment of the idea that the value represents the utility to a player of participating in a game.

Existence of the value is no problem. The value exists in all finite games, and in large classes of NTU games and infinite games. J. F. Mertens (1987) has shown that all non-atomic TU market games have a value. Uniqueness is slightly more delicate. The value is unique in all finite TU games, in many infinite TU games, and in two-person NTU games; in NTU games with three or more players, however, the value is

not always unique. Indeed, we have seen that, in non-atomic NTU markets that are sufficiently smooth, the value coincides with the competitive outcomes, and these are known not to be unique.

17 The Shapley Value: Summary and Conclusions

The Shapley value is perhaps the most “successful” of all cooperative solution concepts. Like the N – M stable sets, it has a very broad spectrum—it almost always exists, and it applies to political as well as economic and “mixed” contexts. It very often gives results with significant intuitive content; unlike with the core, this intuitive content is often of independent interest, not connected with the idea of value in any obvious or transparent manner. Unlike with either the N – M stable sets or the core, there are very general existence theorems, which cover essentially all the applications that one might want to consider.

A very important point is that the value is mathematically tractable. It lends itself to the application of mathematical methods from probability, measure theory, functional analysis and other areas. As a result, a very considerable body of theory has been built around the value; this theory may well be mathematically the richest and deepest in game theory. Of course this is intellectually pleasing, but that is not where its importance lies. Its importance lies in the fact that this theory enables us to deal with the applications, to attack fairly complex models in a systematic manner and to solve them. For non-atomic games there has even been developed a sort of rough calculus that enables us to calculate values “quickly,” though often the rigorous justification takes longer.

Conceptually, the image of the value that arises from the applications is of an index of strength of a player, based on the strength of the coalitions of which he is a member and of those of which he is not a member. A closely related view of the value is as a group decision or arbitrated outcome. Unlike the Nash equilibrium, the core, or the N – M stable set, the value has little or no stability on its own.

18 General Summary and Conclusion

Game-theoretic solution concepts should be understood in terms of their applications, and should be judged by the quantity and quality of their applications. The solution concepts we have considered all have different kinds of applications, which reflect back on the solution concepts and yield different interpretations of them. In each case, important descriptive

and normative insights result; each of the concepts unifies a different aspect of rationality in interactive decision-making.

References

- Akerlof, G. (1970), "The Market for Lemons: Qualitative Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 89, 488–500.
- Anderson, R. (1978), "An Elementary Core Equivalence Theorem," *Econometrica*, 46, 1483–7.
- Arrow, K. J. (1972a), "Models of Job Discrimination," in A. H. Pascal (ed.), *Racial Discrimination in Economic Life*, D. C. Heath, Lexington, Mass., pp. 8–102.
- Arrow, K. J. (1972b), "Some Mathematical Models of Race in the Labor Market," in A. H. Pascal (ed.), *Racial Discrimination in Economic Life*, D. C. Heath, Lexington, Mass., pp. 187–204.
- Arrow, K. J. (1973a), "The Theory of Discrimination," in O. Ashenfelter and A. Rees (eds), *Discrimination in Labor Markets*, Princeton University Press, pp. 3–33.
- Arrow, K. J. (1973b), "Higher Education as a Filter," *Journal of Public Economics*, 2, 193–216.
- Arrow, K. J. and Debreu, G. (1954), "Existence of an Equilibrium for a Competitive Economy," *Econometrica*, 22, 265–90.
- Arrow, K. J. and Hahn, F. (1971), *General Competitive Analysis*, North Holland, New York.
- Aumann, R. J. (1959), "Acceptable Points in General Cooperative n -Person Games," in R. D. Luce and A. W. Tucker (eds), *Contributions to the Theory of Games IV*, Annals of Mathematics Studies, vol. 40, Princeton University Press, pp. 287–324 [Chapter 20].
- Aumann, R. J. (1964), "Markets with a Continuum of Traders," *Econometrica*, 32, 39–50 [Chapter 46].
- Aumann, R. J. (1973), "Disadvantageous Monopolies," *Journal of Economic Theory*, 6, 1–11 [Chapter 48].
- Aumann, R. J. (1974), "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67–96 [Chapter 31].
- Aumann, R. J. (1975), "Values of Markets with a Continuum of Traders," *Econometrica*, 43, 611–46 [Chapter 51].
- Aumann, R. J. (1976), "Agreeing to Disagree," *Annals of Statistics*, 4, 1236–9 [Chapter 32].
- Aumann, R. J. (1979), "On the Rate of Convergence of the Core," *International Economic Review*, 20, 349–57 [Chapter 50].
- Aumann, R. J. (1985), "On the Non-transferable Utility Value: A Comment on the Roth-Shafer Examples," *Econometrica*, 53, 667–77 [Chapter 61c].
- Aumann, R. J. (1990), "Nash Equilibria Are Not Self-Enforcing," in J. J. Gabszewicz, J-F. Richard, and L. Wolsey (eds), *Economic Decision Making: Games, Econometrics, and Optimisation (Essays in Honor of Jacques Dreze)*, Elsevier, Amsterdam, pp. 201–6 [Chapter 34].
- Aumann, R. J. and Kurz, M. (1977), "Power and Taxes," *Econometrica*, 45, 1137–61 [Chapter 52].
- Aumann, R. J., Kurz, M. and Neyman, A. (1983), "Voting for Public Goods," *Review of Economic Studies*, 50, 677–93 [Chapter 56].
- Aumann, R. J., Kurz, M. and Neyman, A. (1987), "Power and Public Goods," *Journal of Economic Theory*, 42, 108–27 [Chapter 55].
- Aumann, R. J. and Shapley, L. S. (1974), *Values of Non-atomic Games*, Princeton University Press.
- Axelrod, R. (1984), *The Evolution of Cooperation*, Basic Books, New York.

- Bertrand, J. (1883), "Theorie mathematique de la Richesse Sociale," *Journal des Savants*, 48, 499–508.
- Bewley, T. (1973), "Edgeworth's Conjecture," *Econometrica*, 41, 425–52.
- Billera, L. J. (1974), "On Games without Side Payments Arising from a General Class of Markets," *Journal of Mathematical Economics*, 1, 129–39.
- Billera, L. J. and Bixby, R. (1973a), "A Characterization of Pareto Surfaces," *Proceedings of the American Mathematical Society*, 41, 261–7.
- Billera, L. J. and Bixby, R. (1973b), "A Characterization of Polyhedral Market Games," *International Journal of Game Theory*, 2, 253–61.
- Billera, L. J. and Bixby, R. (1974), "Market Representation of N -Person Games," *Bulletin of the American Mathematical Society*, 80, 522–6.
- Billera, L. J., Heath, D. C. and Raanan, J. (1978), "International Telephone Billing Rates—A Novel Application of Non-Atomic Game Theory," *Operations Research*, 26, 956–65.
- Binmore, K. (1987), "Perfect Equilibria in Bargaining Models," in K. Binmore and P. Dasgupta (eds), *The Economics of Bargaining*, Blackwell, Oxford, pp. 77–105.
- Bogardi, I. and Szidarovsky, F. (1978), "Application of Game Theory in Water Management," *Applied Mathematical Modelling*, 1, 11–20.
- Bondareva, O. N. (1962), "Teoriia iadra v igre n lits" (Theory of the core in the n -person game), *Vestnik LGU* (Leningrad State University), 13, 141–2.
- Bondareva, O. N. (1963), "Nekotorye primeneniia metodov linejnogo programmirovaniia k teorii kooperativnykh igr" (Some applications of linear programming methods to the theory of cooperative games), *Problemy Kibernetiki*, 10, 119–39.
- Bott, R. (1953), "Symmetric Solutions to Majority Games," in H. W. Kuhn and A. W. Tucker (eds), *Contributions to the Theory of Games II*, Annals of Mathematics Studies, vol. 28, Princeton University Press, pp. 319–23.
- Brown, D. J. and Robinson, A. (1975), "Non-standard Exchange Economies," *Econometrica*, 43, 41–55.
- Butters, G. (1977), "Equilibrium Distribution of Sales and Advertising Prices," *Review of Economic Studies*, 44, 465–92.
- Cheng, H. C. (1981), "What is the Normal Rate of Convergence of the Core?" (Part I), *Econometrica*, 49, 73–84.
- Cournot, A. A. (1838), *Recherches sur les principes mathematiques de la theorie des richesses*, Hachette, Paris (*Researches into the Mathematical Principles of the Theory of Wealth*, trans. N. T. Bacon, Macmillan, New York, 1897).
- Crawford, V. and Knoer, E. M. (1981), "Job Matching with Heterogeneous Firms and Workers," *Econometrica*, 49, 437–50.
- d'Aspremont, C., Gabszewicz, J. J. and Thisse, J. F. (1979), "On Hotelling's 'Stability in Competition,'" *Econometrica*, 47, 1145–50.
- Davis, M. and Maschler, M. (1965). "The Kernel of a Cooperative Game," *Naval Research Logistics Quarterly*, 12, 223–59.
- Davis, M. and Maschler, M. (1967), "Existence of Stable Payoff Configurations for Cooperative Games," in M. Shubik (ed.), *Essays in Mathematical Economics in Honor of Oskar Morgenstern*, Princeton University Press, pp. 39–52.
- Debreu, G. (1952), "A Social Equilibrium Existence Theorem," *Proceedings of the National Academy of Sciences*, 38, 886–93.
- Debreu, G. (1975), "The Rate of Convergence of the Core of an Economy," *Journal of Mathematical Economics*, 2, 1–7.
- Debreu, G. and Scarf, H. (1963), "A Limit Theorem on the Core of an Economy," *International Economic Review*, 4, 235–46.
- Diamond, P. A. (1971), "A Model of Price Adjustment," *Journal of Economic Theory*, 3, 156–68.
- Dierker, E. and Dierker, H. (1972), "The Local Uniqueness of Equilibria," *Econometrica*, 40, 867–82.

- Dreze, J. H., Gepts, S. and Gabszewicz, J. (1969), "On Cores and Competitive Equilibria," in *La Decision*, CNRS, Paris, pp. 91–114.
- Dubins, L. E. and Freedman, D. A. (1981), "Machiavelli and the Gale–Shapley Algorithms," *American Mathematical Monthly*, 88, 485–94.
- Edgeworth, F. Y. (1881), *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*, London.
- Foley, D. (1970), "Lindahl's Solution and the Core of an Economy with Public Goods," *Econometrica*, 38, 66–72.
- Forges, F. (1985), "Correlated Equilibria in a Class of Repeated Games with Incomplete Information," *International Journal of Game Theory*, 14, 129–49.
- Forges, F. (1986), "Correlated Equilibria in Repeated Games with Lack of Information on One Side: A Model with Verifiable Types," *International Journal of Game Theory*, 15, 65–82.
- Fouraker, L. E. and Siegel, S. (1960), *Bargaining and Group Decision Making: Experiments in Bilateral Monopoly*, McGraw-Hill, New York.
- Friedman, J. W. (1977), *Oligopoly and the Theory of Games*, North-Holland, Amsterdam and New York.
- Gabszewicz, J. J. and Dreze, J. H. (1971), "Syndicates of Traders in an Exchange Economy," in H. W. Kuhn and G. P. Szego (eds), *Differential Games and Related Topics*, North-Holland, Amsterdam, pp. 399–414.
- Gabszewicz, J. J. and Mertens, J. F. (1971), "An Equivalence Theorem for the Core of an Economy Whose Atoms Are Not 'Too Big,'" *Econometrica*, 39, 713–22.
- Gale, D. (1984), "Equilibrium in a Discrete Exchange Economy with Money," *International Journal of Game Theory*, 13, 61–4.
- Gale, D. and Shapley, L. S. (1962), "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 69, 9–15.
- Gardner, R. (1975), "Studies in the Theory of Social Institutions," PhD thesis, Cornell University.
- Gardner, R. (1981), "Wealth and Power in a Collegiate Polity," *Journal of Economic Theory*, 25, 353–66.
- Gardner, R. (1984), "Power and Taxes in a One-Party State: The USSR, 1925–1929," *International Economic Review*, 25, 743–55.
- Gately, D. (1974), "Sharing the Gain from Regional Cooperation: A Game Theoretic Application to Planning Investment in Electric Power," *International Economic Review*, 15, 195–208.
- Griesmer, J. H. and Shubik, M. (1963), "Toward a Study of Bidding Processes, Parts I, II and III," *Naval Research Logistics Quarterly*, 10, 11–21, 151–73, 199–217.
- Grodal, B. (1972), "A Second Remark on the Core of an Atomless Economy," *Econometrica*, 40, 581–4.
- Grodal, B. (1975), "The Rate of Convergence of the Core of a Purely Competitive Sequence of Economies," *Journal of Mathematical Economics*, 2, 171–86.
- Guesnerie, R. (1977), "Monopoly, Syndicate and Shapley Value: About Some Conjectures," *Journal of Economic Theory*, 15, 235–52.
- Hammond, P. (1975), "Charity: Altruism or Cooperative Egoism," in E. Phelps (ed.), *Altruism, Morality and Economic Theory*, Russell Sage Foundation, New York.
- Harsanyi, J. C. (1959), "A Bargaining Model for the Cooperative N -Person Game," in A. W. Tucker and R. D. Luce (eds), *Contributions to the Theory of Games IV*, Annals of Mathematics Studies, vol. 40, Princeton University Press, pp. 325–55.
- Harsanyi, J. C. (1967), "Games with Incomplete Information Played by Bayesian Players, Parts I, II and III," *Management Science*, 14, 159–82, 320–34, 486–502.
- Harsanyi, J. C. (1973), "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points," *International Journal of Game Theory*, 2, 1–23.

- Harsanyi, J. C. (1974), "A Solution Concept for N -Person Noncooperative Games," *International Journal of Game Theory*, 3, 211–25.
- Harsanyi, J. C. (1975), "The Tracing Procedure, A Bayesian Approach to Defining a Solution to N -Person Noncooperative Games," *International Journal of Game Theory*, 4, 61–94.
- Hart, S. (1973), "Symmetric Solutions of Some Productive Economies," *International Journal of Game Theory*, 2, 53–62.
- Hart, S. (1974), "Formation of Cartels in Large Markets," *Journal of Economic Theory*, 7, 453–66.
- Hart, S. (1977a), "Values of Non-differentiable Markets with a Continuum of Traders," *Journal of Mathematical Economics*, 4, 103–16.
- Hart, S. (1977b), "Asymptotic Value of Games with a Continuum of Players," *Journal of Mathematical Economics*, 4, 57–80.
- Hart, S. (1980), "Measure-based Values of Market Games," *Mathematics of Operations Research*, 5, 197–228.
- Hart, S. (1985), "Non-zero-sum Two-person Games with Incomplete Information," *Mathematics of Operations Research*, 10, 117–53.
- Hildenbrand, W. (1974), *Core and Equilibria of a Large Economy*, Princeton University Press.
- Hotelling, H. (1929), "Stability in Competition," *Economic Journal*, 39, 41–57.
- Hurwicz, L. (1979a), "Outcome Functions Yielding Walrasian and Lindahl Allocations at Nash Equilibrium Points," *Review of Economic Studies*, 46, 217–25.
- Hurwicz, L. (1979b), "On Allocations Attainable through Nash Equilibria," *Journal of Economic Theory*, 21, 140–65.
- Imai, H. (1983), "On Harsanyi's Solution," *International Journal of Game Theory*, 12, 161–79.
- Kalai, E. and Samet, D. (1984), "Persistent Equilibria in Strategic Games," *International Journal of Game Theory*, 13, 129–44.
- Kalai, E. and Samet, D. (1985), "Unanimity Games and Pareto Optimality," *International Journal of Game Theory*, 14, 41–50.
- Kaneko, M. (1976), "On the Core and Competitive Equilibria of a Market with Indivisible Goods," *Naval Research Logistics Quarterly*, 23, 321–37.
- Kannai, Y. (1970), "Continuity Properties of the Core of a Market," *Econometrica*, 38, 791–815.
- Kannai, Y. (1972), "Continuity Properties of the Core of a Market: A Correction," *Econometrica*, 40, 955–8.
- Knuth, D. E. (1976), *Mariages Stables*, Les Presses de l'Universite de Montreal.
- Kohlberg, E. and Mertens, J. F. (1986), "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003–37.
- Kreps, D., Milgrom, P., Roberts, J. and Wilson, R. (1982), "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 27, 245–52.
- Kreps, D. and Wilson, R. (1982a), "Sequential Equilibria," *Econometrica*, 50, 863–94.
- Kreps, D. and Wilson, R. (1982b), "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253–79.
- Kurz, M. (1977), "Distortion of Preferences, Income Distribution and the Case for a Linear Income Tax," *Journal of Economic Theory*, 14, 291–8.
- Kurz, M. (1978), "Altruism as an Outcome of Social Interaction," *American Economic Review*, 68, 216–22.
- Kurz, M. (1980), "Income Distribution and Distortion of Preferences: The l -Commodity Case," *Journal of Economic Theory*, 22, 99–107.
- Lewis, D. K. (1969), *Convention*, Harvard University Press, Cambridge, Mass.

- Littlechild, S. C. and Owen, G. (1973), "A Simple Expression for the Shapley Value in a Special Case," *Management Science*, 20, 370–2.
- Loehman, E. T., Pengry, D. and Whinston, A. B. (1973), "Cost Allocation for Regional Pollution Treatment Systems," in J. R. Conner and E. T. Loehman (eds), *Economics and Decision Making for Environmental Quality*, University of Florida Press, Gainesville.
- Lucas, W. F. (1969), "The Proof that a Game May Not Have a Solution," *Transactions of the American Mathematical Society*, 137, 219–29.
- Lucas, W. F. and Rabie, M. (1982), "Games with No Solutions and Empty Cores," *Mathematics of Operations Research*, 7, 491–500.
- Luce, R. D. and Raiffa, H. (1957), *Games and Decisions*, John Wiley, New York.
- Malouf, M. and Roth, A. (1981), "Disagreement in Bargaining: An Experimental Study," *Journal of Conflict Resolution*, 25, 329–48.
- Manes, R. P. and Verrecchia, R. (1982), "A New Proposal for Setting Intracompany Transfer Prices," *Accounting and Business Research*, 12, 97–104.
- Mas-Colell, A. (1975), "A Further Result on the Representation of Games by Markets," *Journal of Economic Theory*, 10, 117–22.
- Mas-Colell, A. (1977), "Competition and Value Allocations of Large Exchange Economies," *Journal of Economic Theory*, 14, 419–38.
- Mas-Colell, A. (1980), "Remarks on the Game-theoretic Analysis of a Simple Distribution of Surplus Problem," *International Journal of Game Theory*, 9, 125–40.
- Maschler, M., Peleg, B. and Shapley, L. S. (1972), "The Kernel and Bargaining Set for Convex Games," *International Journal of Game Theory*, 1, 73–93.
- Maynard Smith, J. (1982), *Evolution and the Theory of Games*, Cambridge University Press, Cambridge and New York.
- Mertens, J. F. (1987), "The Shapley Value in the Non-Differentiable Case," *International Journal of Game Theory*, 17, 1–65.
- Milgrom, P. and Roberts, J. (1982a), "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis," *Econometrica*, 50, 443–60.
- Milgrom, P. and Roberts, J. (1982b), "Predation, Reputation and Entry Deterrence," *Journal of Economic Theory*, 27, 280–312.
- Milgrom, P. and Weber, R. (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–122.
- Milnor, J. W. and Shapley, L. S. (1978), "Values of Large Games II: Oceanic Games," *Mathematics of Operations Research*, 3, 290–307.
- Myerson, R. (1978), "Refinements of the Nash Equilibrium Concept," *International Journal of Game Theory*, 7, 73–80.
- Myerson, R. (1981), "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.
- Nash, J. F. (1950), "The Bargaining Problem," *Econometrica*, 18, 155–62.
- Nash, J. F. (1951), "Non-cooperative Games," *Annals of Mathematics*, 54, 86–295.
- Nash, J. F. (1953), "Two-person Cooperative Games," *Econometrica*, 21, 128–40.
- Neyman, A. (1985), "Semi-values of Political Economic Games," *Mathematics of Operations Research*, 10, 390–402.
- Novshek, W. (1980), "Cournot Equilibrium with Free Entry," *Review of Economic Studies*, 47, 473–86.
- Novshek, W. and Sonnenschein, H. (1978), "Cournot and Walras Equilibrium," *Journal of Economic Theory*, 19, 223–66.
- Osborne, M. J. (1984), "Why Do Some Goods Bear Higher Taxes Than Others?," *Journal of Economic Theory*, 32, 301–16.
- Ostroy, J. (1981), "Differentiability as Convergence to Perfectly Competitive Equilibrium," *Journal of Mathematical Economics*, 8, 59–73.

- Parker, T. (1943), "Allocation of the Tennessee Valley Authority Projects," *Transactions of the American Society of Civil Engineering*, 108, 174–87.
- Pauly, M. V. and Satterthwaite, M. A. (1981), "The Pricing of Primary Care Physicians' Services: A Test of the Role of Consumer Information," *Bell Journal of Economics*, 12, 488–506.
- Peleg, B. (1967), "Existence Theorem for the Bargaining Set $M_1^{(i)}$," in M. Shubik (ed.), *Essays in Mathematical Economics in Honor of Oskar Morgenstern*, Princeton University Press, pp. 53–6.
- Prescott, E. C. and Visscher, M. (1977), "Sequential Location among Firms with Foresight," *Bell Journal of Economics*, 8, 378–93.
- Quinzii, M. (1984), "Core and Competitive Equilibria with Indivisibilities," *International Journal of Game Theory*, 13, 41–60.
- Radner, R. (1980), "Collusive Behavior in Noncooperative Epsilon-equilibria of Oligopolies with Long but Finite Lives," *Journal of Economic Theory*, 22, 136–54.
- Radner, R. (1981), "Monitoring Cooperative Agreements in a Repeated Principal–Agent Relationship," *Econometrica*, 49, 1127–48.
- Ransmeier, J. S. (1942), *The Tennessee Valley Authority: A Case Study in the Economics of Multiple Purpose Stream Planning*, Vanderbilt University Press, Nashville, Tenn.
- Riker, W. and Shapley, L. S. (1968), "Weighted Voting: A Mathematical Analysis for Instrumental Judgments," in J. R. Pennock and J. W. Chapman (eds), *Representation: Nomos X*, Atherton Press, New York, pp. 199–216.
- Rosenmuller, J. (1981), "Values of Non-sidepayment Games and Their Application in the Theory of Public Goods," in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, vol. 4 of *Gesellschaft, Recht, Wirtschaft*, Bibliographisches Institut, Mannheim, Vienna and Zurich, pp. 111–29.
- Rosenmuller, J. (1982), "On Values, Location Conflicts, and Public Goods," in M. Deistler, E. Furst and G. Schwodiauer (eds), *Proceedings of the Oskar Morgenstern Symposium at Vienna*, Physica-Verlag, Vienna, pp. 74–100.
- Rosenthal, R. W. (1979), "Sequences of Games with Varying Opponents," *Econometrica*, 47, 1353–66.
- Ross, S. (1973), "The Economic Theory of Agency: The Principal's Problem," *American Economic Review*, 63, 134–9.
- Roth, A. E. (1977), "The Shapley Value as a von Neumann–Morgenstern Utility," *Econometrica*, 45, 657–64.
- Roth, A. E. (1982), "The Economics of Matching: Stability and Incentives," *Mathematics of Operations Research*, 7, 617–28.
- Roth, A. E. (1984), "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory," *Journal of Political Economy*, 92, 991–1016.
- Roth, A. E., Malouf, M. and Murnighan, J. K. (1981), "Sociological versus Strategic Factors in Bargaining," *Journal of Economic Behavior and Organization*, 2, 153–77.
- Roth, A. E. and Murnighan, J. K. (1982), "The Role of Information in Bargaining: An Experimental Study," *Econometrica*, 50, 1123–42.
- Rothschild, M. and Stiglitz, J. (1976), "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 90, 629–49.
- Rubinstein, A. (1979), "An Optimal Conviction Policy for Offenses that May Have Been Committed by Accident," in S. J. Brams, A. Schotter and G. Schwodiauer (eds), *Applied Game Theory*, Physica-Verlag, Vienna, pp. 406–13.
- Rubinstein, A. (1982), "Perfect Equilibria in a Bargaining Model," *Econometrica*, 50, 97–110.
- Rubinstein, A. (1995), "Equilibrium in Supergames," in N. Megiddo (ed.), *Essays in Game Theory in Honor of Michael Maschler*, Springer-Verlag, Berlin, pp. 17–27.
- Salop, S. C. (1979), "Monopolistic Competition with Outside Goods," *Bell Journal of Economics*, 10, 141–56.

- Scarf, H. E. (1967), "The Core of an N -Person Game," *Econometrica*, 35, 50–69.
- Scarf, H. E. (1973), *The Computation of Economic Equilibria*, Yale University Press, New Haven, Conn.
- Schmeidler, D. (1969a), "Competitive Equilibria in Markets with a Continuum of Traders and Incomplete Preferences," *Econometrica*, 37, 578–85.
- Schmeidler, D. (1969b), "The Nucleolus of a Characteristic Function Game," *Society for Industrial and Applied Mathematics, Journal on Applied Mathematics*, 17, 1163–70.
- Schmeidler, D. (1972), "A Remark on the Core of an Atomless Economy," *Econometrica*, 40, 579–80.
- Schmeidler, D. (1980), "Walrasian Analysis via Strategic Outcome Functions," *Econometrica*, 48, 1585–94.
- Scott, M. (1981), "Comments on R. E. Verrecchia's 'A Question of Equity: Use of the Shapley Value to Allocate State and Local Income and Franchise Taxes'", in S. Moriarity (ed.), *Joint Cost Allocations: Proceedings of the University of Oklahoma Conference on Cost Allocation, April 23, 24, 1981*, Center of Economic and Management Research, University of Oklahoma.
- Selten, R. (1965), "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrage-
tragheit," *Zeitschrift für die gesamte Staatswissenschaft*, 121, 301–24, 667–89.
- Selten, R. (1973), "A Simple Model of Imperfect Competition where 4 are Few and 6 are Many," *International Journal of Game Theory*, 2, 141–201.
- Selten, R. (1975), "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25–55.
- Selten, R. (1980), "A Note on Evolutionarily Stable Strategies in Asymmetric Animal Conflicts," *Journal of Theoretical Biology*, 84, 93–101.
- Selten, R. and Güth, W. (1982), "Game Theoretic Analysis of Wage Bargaining in a Simple Business Cycle Model," *Journal of Mathematical Economics*, 10, 177–95.
- Shapiro, N. Z. and Shapley, L. S. (1978), "Values of Large Games I: A Limit Theorem," *Mathematics of Operations Research*, 3, 1–9.
- Shapley, L. S. (1953), "A Value for N -Person Games," in H. W. Kuhn and A. W. Tucker (eds), *Contributions to the Theory of Games II*, Annals of Mathematics Studies, vol. 28, Princeton University Press, pp. 307–12.
- Shapley, L. S. (1967), "On Balanced Sets and Cores," *Naval Research Logistics Quarterly*, 14, 453–60.
- Shapley, L. S. (1969), "Utility Comparison and the Theory of Games," in *La Decision*, CNRS, Paris, pp. 251–63.
- Shapley, L. S. (1971), "Cores of Convex Games," *International Journal of Game Theory*, 1, 11–26.
- Shapley, L. S. (1975), "An Example of a Slow Converging Core," *International Economic Review*, 16, 345–51.
- Shapley, L. S. and Scarf, H. E. (1974), "On Cores and Indivisibilities," *Journal of Mathematical Economics*, 1, 23–8.
- Shapley, L. S. and Shubik, M. (1954), "A Method for Evaluating the Distribution of Power in a Committee System," *American Political Science Review*, 48, 787–92.
- Shapley, L. S. and Shubik, M. (1967), "Concepts and Theories of Pure Competition," in M. Shubik (ed.), *Essays in Mathematical Economics in Honor of Oskar Morgenstern*, Princeton University Press.
- Shapley, L. S. and Shubik, M. (1969), "On Market Games," *Journal of Economic Theory*, 1, 9–25.
- Shapley, L. S. and Shubik, M. (1972), "The Assignment Game I: The Core," *International Journal of Game Theory*, 1, 111–30.
- Sharkey, W. W. (1982), "Suggestions for a Game Theoretic Approach to Public Utility Pricing and Cost Allocation," *Bell Journal of Economics*, 13, 57–68.

- Shavell, S. (1979), "Risk Sharing and Incentives in the Principal and Agent Relationship," *Bell Journal of Economics*, 10, 55–73.
- Shitovitz, B. (1973), "Oligopoly in Markets with a Continuum of Traders," *Econometrica*, 41, 467–502.
- Shitovitz, B. (1974), "On Some Problems Arising in Markets with Some Large Traders and a Continuum of Small Traders," *Journal of Economic Theory*, 8, 458–70.
- Shubik, M. (1959), "Edgeworth Market Games," in R. D. Luce and A. W. Tucker (eds), *Contributions to the Theory of Games IV*, Annals of Mathematics Studies, vol. 40, Princeton University Press, pp. 267–78.
- Shubik, M. (1963), "Incentives, Decentralized Control, the Assignment of Joint Costs and Internal Pricing," in C. Bonini, R. Jaedicke and M. Wagner (eds), *Management Controls*, McGraw-Hill, New York.
- Sondermann, D. (1974), "Economies of Scale and Equilibria in Coalition Production Economies," *Journal of Economic Theory* 8, 259–91.
- Spence, M. (1974), *Market Signaling: Information Transfer in Hiring and Related Processes*, Harvard University Press, Cambridge, Mass.
- Spence, M. (1976), "Competition in Salaries and Signaling Prerequisites for Jobs," *Quarterly Journal of Economics*, 90, 51–75.
- Stearns, R. (1965), "A Game without Side Payments that has No Solution," in W. F. Lucas (ed.), Report of the Fifth Conference on Game Theory, mimeograph, Department of Mathematics, Princeton University.
- Straffin, P. D. and Heaney, J. P. (1981), "Game Theory and the Tennessee Valley Authority," *International Journal of Game Theory*, 10, 35–43.
- Suzuki, M. and Nakayama, M. (1976), "The Cost Assignment of Cooperative Water Resource Development: A Game Theoretic Approach," *Management Science*, 22, 1081–6.
- Thompson, G. L. (1980), "Computing the Core of a Market Game," in A. V. Fiacco and K. O. Kortanek (eds), *External Methods and System Analysis*, Springer-Verlag, Vienna.
- Verrecchia, R. E. (1981), "A Question of Equity: Use of the Shapley Value to Allocate State and Local Income and Franchise Taxes," in S. Moriarity (ed.), *Joint Cost Allocations: Proceedings of the University of Oklahoma Conference on Cost Allocation, April 23, 24, 1981*, Center for Economic and Management Research, University of Oklahoma.
- Vickrey, W. (1961), "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, 16, 8–37.
- Vickrey, W. (1962), "Auctions and Bidding Games," in *Recent Advances in Game Theory* (Conference Proceedings), Princeton University Press, pp. 15–27.
- Vind, K. (1964), "Edgeworth Allocations in an Exchange Economy with Many Traders," *International Economic Review*, 5, 165–77.
- Vind, K. (1972), "A Third Remark on the Core of an Atomless Economy," *Econometrica*, 40, 585–6.
- von Neumann, J. (1937), "A Model of General Economic Equilibrium," *Review of Economic Studies*, 13, 1–9.
- von Neumann, J. and Morgenstern, O. (1944), *Theory of Games and Economic Behavior*, Princeton University Press.
- Wilson, R. (1979), "Auctions of Shares," *Quarterly Journal of Economics*, 93, 675–98.

Note added in proof Professor Elmer Offenbacher of Temple University has informed us that Newtonian mechanics, too, predicts the precession of the perihelion of Mercury (see Section 2)—that relativity theory only provides a correction in the *magnitude* of the effect. Indeed, Newtonian mechanics predicts a precession of 432 arc seconds per century, whereas the observed value—which is also that predicted by relativity—is 475.