4 CORE as a Macrocosm of Game-Theoretic Research, 1967–1987

It's really a tremendous pleasure to participate in this birthday celebration. My first visit to CORE was in 1969, and I was immediately captivated by the dynamism of the place and the people who constitute it. Since then there have been many visits, and I have come to feel about CORE as my European home.

When Henry Tulkens asked me for a title for this talk, I first thought of something like "CORE as a Center for Game-Theoretic Research." But after all, CORE has been much more than such a center. It is not only that important game-theoretic research has been going on here at CORE for the last twenty years. It is that most of the central, most interesting, and most vital strands of game theory have been represented here. So I considered "CORE as a Microcosm of Game Theory." But this doesn't do justice to the subject either. "Microcosm" means "small world." For game theory, CORE has been not a small world but a big world. In the last ten or fifteen years the influence of game theory on economics has been growing steadily. CORE has played a very important part in this development, both in creating the game theory itself and in working out its applications to economics. CORE looms very large in the world of game theory. It is a macrocosm, not a microcosm.

Let me mention here just some of the important areas in game theory to which CORE has made major, central contributions:

1. Economies with large indivisible units (syndicates) (Drèze, Gabszewicz, Gepts, Mertens; Shitovitz, Greenberg, Postlewaite)
2. Economies with many goods (Mertens, Gabszewicz; Bewley)
3. Core of an economy (Hildenbrand, Bohm, Champsaur, Gepts; Schmeidler, Zamir, Bergstrom, Roberts, Vind)
4. Repeated and stochastic games (Mertens, Sorin, Forges, Waternaux, Lefevre; Zamir, Neyman, Hart, Ponssard)
5. Shapley value, both abstract theory and applications to economics (Mertens, Drèze; Neyman, Tauman, Mirman, Aumann)
6. Strategic equilibria (Moulin, Vial, Mertens; Kohlberg, Schmeidler, Aumann)
7. Coalitions, coalitional games, and coalition structures (Drèze, Delbaen, Tulkens, d'Aspremont; Sondermann, Zamir, Schmeidler, Greenberg, Aumann)

This chapter originally appeared in *Contributions to Operations Research and Economics: The Twentieth Anniversary of CORE*, edited by B. Cornet and H. Tulkens, pp. 5–16, The MIT Press, Cambridge and London, 1989.

8. Incentives, incomplete information, mechanisms (d'Aspremont, Gerard-Varet, Schoumaker)

9. Oligopoly (Gabszewicz, Deschamps, Jacquemin, Vial, d'Aspremont; Hansen, Roberts, Sonnenschein)

10. Public goods (Drèze, de la Vallée Poussin, Champsaur, Tulkens; Henry, Roberts, Rosenthal, Weber, Zamir)

This is a very impressive list. The centrality of CORE's work in these areas is difficult to overestimate. If one wants to know what the classic contributions are, it is enough to look at the list of CORE reprints. The seminal contribution that started area 1—economies with atoms—is the classic paper of Drèze, Jaskold-Gabszewicz, and Gepts; and since then most of the work in this area has been done right here at CORE, both by local researchers and by visitors. Similar statements can be made for the theory of repeated games, Shapley value, Core, and many of the other areas mentioned above. I vividly remember the excitement in the air at the time of my first visit, on the subject of Cores and equilibria of economies with a continuum of goods *and* a continuum of agents. Does the equivalence principle still hold? The question is still very much with us. That work in the late 1960s was the forerunner of a stream of research that has continued to occupy researchers to this day.

I could go on reminiscing in this way. (It is, I suppose, a sign of age; CORE may be a vigorous twenty, but some of the rest of us are a little older.) Instead, let me speculate a little on what it is that makes CORE so unique. There are, after all, several great centers in the world in the areas of mathematical economics and game theory, staffed by extremely able, original, friendly, and cooperative people. But at CORE there is a quality of intensity and excitement that is difficult to find elsewhere. Why?

The reason appears to be a very unusual mix of outstanding permanent staff and visitors. Many institutions have excellent permanent staff, as well as a budget for a few visitors, and perhaps an occasional emphasis year. Other institutions are built primarily on visitors. CORE, with its strong permanent core plus a considerable stream of visitors, provides a unique breeding ground: a place where cross-fertilization leads to the conception of new ideas, as well as a womb—a warm, supportive environment in which these ideas can grow and mature.

For what it may be worth, much of my own best work is inextricably tied up with this place. Perhaps an appropriate modesty would have called for the omission of my name from the above summary. But I am both grateful and proud to have been associated with CORE, and to continue to be; I have no desire to hide the association, or to make light of it.

Ladies and gentlemen, we must get to some substance. The most substantial evidence of CORE's success is in its accomplishments. One cannot hope to survey game theory at CORE in one lecture. What I propose to do here is to take two or three ideas that were generated at CORE and see where they have led on the intuitive level.

One of the areas in which CORE has made very important, central contributions is the theory and applications of the Shapley value. The Shapley value is a kind of index of the power of each participant in an interactive situation—each player in a game. It is to a game what the mean is to a distribution; without making any specific predictions, it gives one an idea of how the land lies. Since 1974 much of the most important theoretical work on values of large games has been done at CORE; worthy of particular mention in this connection is a very important piece that appeared in the *International Journal of Game Theory* in early 1988, in which Jean-François Mertens solves a difficult problem that had been open for many years—the existence of a value on an extremely large class of games, including all market games, differentiable or not. In the applications, too, CORE has been very active. Jacques Drèze and I spent seven years working on the value of fixed price markets; the results appeared in the November 1986 issue of *Econometrica*—just in time for CORE's birthday. More about that later. First I would like to discuss an application of the value to public goods, which is another subject on which CORE has been very active. This represents joint work of M. Kurz and A. Neyman with me, which appeared in the *Journal of Economic Theory* in June 1987. The particular aspects that I will discuss now were mostly generated at the time of a short visit of Mordecai Kurz to CORE in 1978, when I was spending half a year here.

The situation we were considering is that of a country in which the choice of public goods is determined by a majority vote of the citizens. There may, however, be residents of the country who are not citizens; they pay taxes but are not allowed to vote. The public goods we are talking about are nonexclusive. Once produced, they may be used by anybody, citizens as well as noncitizens, and, in the case of citizens, quite independently of how they voted.

In this kind of situation it seems obvious that it is very important what the tastes of the citizens are, and that the tastes of the noncitizens count for nothing. But value analysis leads to precisely the opposite conclusion

—that it does not matter who the citizens are. For a given distribution of tastes of all residents, value analysis leads to the same outcome no matter what tastes the citizens have.

To bring this point home, let me illustrate with an example. There are two kinds of agents in the country: one kind prefers public libraries; the other prefers television. We can think of the country as equally divided on this score. Then the value of a library lover—his power, so to speak—is not only independent of whether he is a citizen but also of whether there is any library lover at all who is a citizen. The value outcome is the same whether all the citizens are library lovers or all are television fans.

Now this is totally unexpected, it sounds absurd. But the value has a very good track record. There have been several instances in which the value seemed strange at first but very few, if any, where this strangeness survived closer examination. Indeed, if theory always yields expected results, it is neither very interesting nor very important. The best kind of theory is that which sounds strange at first but on closer examination turns out to be "right" after all, to yield an important *unexpected* insight.

When we look at the voting situation more carefully, the result under discussion makes perfectly good economic sense. In a cooperative game people may be bribed to vote in one way or another. The individual voter knows that no matter how he votes, he will be able to use the public goods that will in fact be produced. He knows that his vote may influence the decision as to which public goods will be produced, but he also knows that this influence is quite small. For the right price, he might be willing to sell his vote.

What is the minimal price at which a citizen will sell his vote? In any case it has to be more than the expected result of his voting in the way he wants. If there are ten million citizens in the country, his expected influence on the voting process is, for him, worth about one ten-millionth of *his* expected usage of the public goods; if he is a television fan, say, this is about one ten-millionth of *his* television viewing, which is a quite negligible amount. So he should be willing to sell his vote for next to nothing.

What will the resulting market price of a vote be? Well, one needs only 51 percent of the voters to put through any public goods program one wants. It doesn't take very long to realize that in such a situation, cutthroat competition will develop among the voters, and the price of a vote will go down to practically nothing. So the tastes of the *voters* (as distinguished from plain residents) don't make a difference.

If the vote doesn't determine the shape of the public goods program, what does? It is the taxes paid by the taxpayer; it is who *pays* for the public goods, not who votes for them. All this is set out in the above-mentioned JET paper.

Very good. So the value of a vote makes economic sense after all. But it probably sounds quite strange to the man in the street. "If that is economic sense," he will say, "so much the worse for economics."

The cynics among us will say that this just confirms what they have always said, that democracy is a sham, that all that counts is money. But how do the rest of us make common sense, not just "economic sense," of the result? Can the vote really count for next to nothing in determining the mix of public goods?

The explanation of the paradox lies in the underlying assumption that a vote can be sold. I am not talking about legal, "administrative" restrictions; those can always be overcome in some way, and usually they are. I am talking about something that is more effective, namely, informational restrictions. In most countries the vote is secret. You can't sell a vote because there is no way that you can guarantee delivery. So the result collapses. It's not a cooperative game; the contracts are unenforceable (recall that a *cooperative* game is one in which agreements are enforceable). The economic argument behind the result also collapses.

Does that make the result useless, a waste of time? Not at all; on the contrary, it becomes more beautiful. It has always been assumed that the reason for the secret vote is to protect the voter from the pressure of others to influence his vote. But it turns out that there is a more intrinsic reason: to protect the voter from himself, or rather to protect the voters from themselves. If we would enable a vote to be sold (note that I say "enable," not "allow"), all voters would want to sell their votes, and then because of the competition effect, they would get next to nothing for them.

This line of reasoning also has some implications for the rules of voting in a legislature. Most voting in a legislature is required to be open. Is this good or bad? Well, it would seem to be bad, because it does enable a vote to be sold, not necessarily for money but for other votes. This is what makes logrolling possible. But logrolling is not altogether a bad thing; it enables the legislators (and indirectly, the voters) to express the strengths of their preferences and have them taken into account. In some sense the open vote is a more perfect expression of democracy just *because* it blunts the tyranny of the majority.

The bottom line is that game theory cannot tell you what kind of social effects to seek. But it *can* tell you what kind of effects to expect from the social system that you build, and sometimes they are far from obvious.

The next idea I will discuss is taken from the theory of repeated games with incomplete information. Repeated games are game theory's tool for dealing with informational effects over time: teaching and learning, reputations, mutual help, establishing cooperative or inimical relationships, signaling, revealing or concealing substantive information, and so on.

One of the earliest insights yielded by the theory was that in the case of incomplete information—when one side has substantive information that the other side does not—the informed side cannot make use of its special information without eventually revealing it. The revelation is indirect; if one side makes use of its information, the other side will eventually infer the information from the actions of the first side.

This is a gross effect; it is true "in the large." But there are some very interesting second-order effects, which I would like to discuss now, that indicate that "in the small" perhaps the informed side *can* make use of its information. The work I am talking about was started by Jean-François Mertens and Shmuel Zamir in a 1976 paper published in the *International Journal of Game Theory* (CORE reprint 312), and I understand that it has recently been developed further. The optimal strategy when you don't want to reveal your information often calls for randomization, "confusing the enemy" by sometimes doing one thing, sometimes the other. The way in which one randomizes—the relative frequency with which one takes the one action or the other—must not depend on the substantive information that one is trying to conceal, for otherwise, the other side could read the information from the relative frequency. But when one randomizes, there is a natural random variation in the series of outcomes. Could the informed side "hide" behind this natural random variation and get a result that is a little better for it than could be obtained if it made no use of its information?

A good parable for this situation is a story I heard many years ago while still in high school. By law, the bread that bakeries sell must weigh what it is claimed to weigh; a 1-kilo loaf must weigh a kilo, and so on. But in baking bread, it is difficult to be that precise. So the law permits some leeway: the standard deviation of weight that might occur in the normal course of events is estimated, and a shortfall is allowed by that amount.

A customer of a certain bakery noticed that he was consistently getting bread that was slightly underweight. The baker seemed to be within his rights because the shortfall was within the legal limits. So what the customer did was to keep a record of the weights of the loaves he bought. Although each one was within the legal limit, the record did not have the distribution it should have had (approximately normal). With this information the customer went to the police, and because of the way the law was worded, the police were eventually able to obtain a conviction of the baker for deliberately selling underweight bread.

This parable seems to indicate that the informed player cannot hide behind randomness in order to take advantage of his information. Curiously, that is not always the case. There are games in which the informed player can take advantage of his special information, up to the order of magnitude of one standard deviation (approximately the square root of the number of repetitions). The reason is that though the uninformed player will eventually "catch on," by the time he does, it is too late to do anything about it. These are fairly "flat" games in which revealing information will not significantly harm the informed player. In other games, which are not "flat," the extent to which the informed player can use his information to get a second-order advantage, while still concealing it, is much smaller.

We return finally to the subject of values of fixed price economies, to which I alluded earlier. As we all know, Jacques Drèze has made fundamental contributions to this subject, including what is known as the "Drèze equilibrium" for such markets. I have done some work on the theory and applications of the Shapley value. So when I arrived at Louvain-la-Neuve in the winter of 1978 to spend a semester at CORE, Jacques and I decided to join forces to have a look at what the Shapley value has to say about fixed price markets. This led to seven years of some of the most intense, frustrating, and rewarding work that I have ever done, and I am confident that for Jacques it was much the same.

Let me try to give you just a flavor of our investigation. It has been known for a long time that in ordinary economies that are "free" (not constrained by fixed prices) and competitive (have many individually insignificant agents), the Shapley value leads to the competitive equilibrium. In other words, the a priori power of each agent in such an economy is precisely what is dictated by the law of supply and demand; what an agent can "expect" to get in the game of free, unconstrained exchange is whatever his net worth will buy him, at prices determined so that total supply matches total demand.

Fixing prices means constraining the set of trades—and therefore of consumptions—available to each agent. An agent can now consume only what he can buy (at the fixed prices) with the worth (at the fixed prices) of his endowment. Thus a fixed price market looks much like an ordinary market—not the originally given market, but one derived from it by replacing each consumption set by the set of all consumptions that are attainable from the endowment by trading at the fixed prices.

So, calculating Shapley values for the original market, with fixed prices, should be equivalent to calculating Shapley values for the derived market, but without fixing prices in the derived market. By the principle enunciated above, one would expect to be led to the competitive equilibria of the derived market. These equilibria of course have nothing to do with competitive equilibria in the original market, and the prices arising endogenously in the derived market have nothing to do with prices in the original market. They are in a different currency, somewhat like the rationing coupons that governments issue to achieve a "fair" or at least orderly distribution of scarce commodities in times of crisis (such as war).

Much of this had been realized already by Drèze and Müller ("Optimality Properties of Rationing Schemes," *Journal of Economic Theory*, 1980), who gave the name "coupons" to the currency that arises endogenously in the derived market. But their work was not in the context of Shapley values.

The analogy with rationing coupons is particularly apt in several ways. Rationing coupons do not replace money, they come in addition to money; purchases are paid for in both ordinary money *and* coupons. Moreover the shortages that necessitate rationing are impossible in a "free" market; the prices of scarce commodities are simply driven so high that there is no excess demand for them. Shortages (and surpluses) are *created* by fixed prices; since the prices are not set to equate supply and demand, excess demands and supplies ensue, and this is what leads to rationing.

In other ways, though, the analogy is less apt. Rationing coupons are usually issued separately for each of many scarce commodities (oil, sugar, gasoline, etc.); here there is only one kind of coupon, which is used for all commodities. (The usual kind of rationing, in which different goods are rationed separately, is what lies at the basis of the "Drèze equilibrium" for fixed price markets.) Moreover coupon "prices" (the coupon amount to be paid per unit of each commodity) are normally fixed themselves; they are set exogenously. By contrast, in the equilibrium discussed above (to which the Shapley value appears to lead), the coupon prices are endogenous; they are set by market forces. Finally, in the usual kind of rationing, each citizen is granted some positive amount of coupons to start with, a "coupon endowment." Here, however, this seems not to be the case, since in equilibrium, each citizen may spend in coupons at most what he receives from the goods that he sells.

The problem with this analysis was that one of the classical Arrow-Debreu conditions that guarantees the existence of a competitive equilibrium—the nonsatiation condition—fails for the derived market, and indeed most of these markets do not have any competitive equilibria. This had already been realized by Drèze and Müller. In the context of Shapley values, this seemed at the time very odd because we did not see how the association of values with competitive equilibria depends on nonsatiation, and we *were* able to prove that a value exists (as it does in just about all applications, which is one of its advantages). This was only the first in a series of "contradictions in mathematics" that appeared to block our path, and it was a long time before we began to see, even dimly, what is going on with the values of these markets.

The whole story, with all its twists and turns, is too long to be told here. Here is the conclusion we eventually reached: When the derived market does not have a competitive equilibrium, one gives each trader a positive coupon endowment after all; his budget then consists of this coupon endowment *plus* the coupon proceeds from the sale of his goods endowment, and he may buy what he wants with this budget. (One must remember that purchases must be "legal": They must be paid for in money, at the fixed prices, as well as in coupons.) For appropriately chosen coupon endowments and coupon prices, supply then matches demand. Each Shapley value of the original fixed price market corresponds to some such "equilibrium."

This may not look like much at first. The Shapley value is by definition Pareto optimal (given that one must trade at the fixed money prices), and the second law of welfare economics implies that *every* such outcome can be "supported" by appropriate coupon prices and endowments. But in general, the coupon endowments of some agents may be negative, of others zero, and of still others positive. What is asserted here is that when the derived market has no competitive equilibrium, each Shapley value can be supported by coupon endowments that are *positive* for all traders.

Even more can be said about the coupon endowments corresponding to Shapley values. Most important, the (endogenous) coupon endowment of each trader depends only on the (exogenous) goods endowment of that trader, not on his tastes (utility function). Two traders with the same goods endowment but different tastes will get the same coupon endowment. And the dependence is monotonic: If trader Adams starts out with at least as much of each good as trader Brown, then Adams's coupon endowment will be at least as large as Brown's.

For example, fixed wages may lead to an oversupply of labor. To correct this, various rationing schemes that involve cutting down on working hours have been proposed. In the scheme implied by the Shapley value, the maximum work week for any worker would depend on his time endowment —on how much time he has. A worker who for some reason can only work part-time might be assigned a quota smaller than the average, even though he may be able to fill the average quota.

Perhaps the most important conclusion to be drawn from all this is the following: Although the Shapley values of fixed price markets are not generally associated with competitive equilibria, they have one important property in common with competitive equilibria. At a competitive equilibrium each trader is assigned a choice set—the budget set—from which he chooses an element that he most prefers. Of course the element chosen by the trader depends on his tastes; traders with the same budget set but with different utilities will in general wind up with different consumptions. But the *choice set* itself does not depend on the trader's tastes; it depends only on his endowment. I mean by this that agents with the same endowments have the same choice set, irrespective of their utilities. Of course budget sets depend on prices, and in a sense prices "depend" on everyone's utilities. Nevertheless, once determined, they are the same for everyone.

It is this basic property that is shared by the Shapley values of fixed price markets. Each trader is assigned a choice set, namely, the "coupon budget set," defined as the set of all consumptions that the trader can legally purchase with the sum of his coupon endowment and the proceeds of the sale of his goods endowment at the going coupon prices. From this choice set the trader chooses a bundle that he most prefers according to his tastes. But the choice set itself does not depend on his tastes, it depends on his goods endowment only, in exactly the same sense that in an ordinary market, a trader's budget set depends on his goods endowment only.

Lest you think that this is a rather common phenomenon, let me assure you that it is not. In many situations the Shapley value depends on the *cardinal* utility function of an agent, on his attitude toward risk; two agents who differ only in their attitudes toward risk—have the same ordinal preferences and the same physical and legal characteristics (endowment, right to vote, etc.)—may yet be assigned considerably different outcomes by the Shapley value. This can even happen when the situation under consideration involves no explicit element of risk—when the only elements of risk are "strategic," that is, due to uncertainty as to how other players will behave. When it does happen, the outcome cannot be represented as the result of optimizing over a choice set that is independent of tastes, since such optima would be independent of attitudes toward risk.

Some examples where the Shapley value depends crucially on attitudes toward risk are ordinary markets with finitely many traders, the public goods economies with many agents considered above, and taxation models with many agents (like those that Kurz and I consider in *Econometrica*, 1977). In these and other cases, agents who are "fearful," or risk averse, are often penalized by being assigned smaller choice sets. It is then impossible to "decouple" the Shapley value, to represent it as the outcome of each trader choosing from a choice set that is independent of his tastes. I venture to say that where this kind of decoupling is impossible is the "normal," expected case—that the decoupling occurring in ordinary, or fixed price, markets with many traders constitutes a very special phenomenon.

What is the significance of this phenomenon? It has to do with the ideas of nondiscrimination, of equal opportunity, of anonymity, and these in turn are closely associated with "perfect competition." When you bargain for a purchase, the price may well depend on your tastes or your attitudes toward risk. But the very fact that you are bargaining indicates that the situation is not perfectly competitive. In a perfectly competitive situation, your tastes don't matter; if you don't get the "right" price from one agent, you will get it from another. Discrimination—distinguishing agents by their tastes—is widely considered an anticompetitive practice. Of course agents with different physical endowments should not expect equal opportunities. But the requirement that agents with the same physical endowments be allowed to choose from the same choice set does appear nicely to capture the idea of equal opportunity.

The point is not to give such agents equal opportunity by fiat. Rather, the situation itself should be intrinsically nondiscriminatory: The opportunities for each agent that are implicit in the situation should be independent of his tastes. This is exactly what the above condition on the Shapley value says. We thus conclude that whereas fixing prices does restrict consumption sets, and changes outcomes drastically, it does not change the inherently competitive nature of the market.

One often thinks of perfect competition in terms of large numbers of small agents. This condition complements the above nicely; without it, we are unlikely to get nondiscrimination. But the presence of many small agents is not sufficient by itself to ensure perfect competition; it does not rule out imperfections like public goods, increasing returns, transactions costs, and so on. Moreover, it is perhaps better to think of many small agents as a physical condition leading to perfect competition, whereas the idea of intrinsic nondiscrimination is more like a definition, it is close to what we *mean* by perfect competition.

The thrust of my remarks has been to give you some examples of how game-theoretic work at CORE has yielded practical and theoretical insights of a general nature that are expressible in words rather than formulas. That is about all that game theory—or for that matter, economic theory—is good for anyway: qualitative insights. I don't think we are going to get much that is quantitatively useful from our disciplines.

Let me close with a word of appreciation and gratitude that transcends the personal. CORE has been very hospitable not only to me but to a large number of Israeli game theorists. For us, the relationship has been enormously fruitful. I can only hope that it has not all been one way.

Thank you.