

Click to prove
you're human



First published Sat Jan 25, 1997; substantive revision Sun Sep 3, 2023

Game theory in the form known to economists, social scientists, and biologists, was given its first general mathematical formulation by John von Neumann and Oskar Morgenstern (1944). For reasons to be discussed later, limitations in their formal framework initially made the theory applicable only under special and limited conditions. This situation has dramatically changed, in ways we will examine as we go along, over the past seven decades, as the framework has been deepened and generalized. Refinements are still being made, and we will review a few outstanding problems that lie along the advancing front edge of these developments towards the end of the article. However, since at least the late 1970s it has been possible to say with confidence that game theory is the most important and useful tool in the analyst's kit whenever she confronts situations in which what counts as one agent's best action (for her) depends on expectations about what one or more other agents will do, and what counts as their best actions (for them) similarly depend on expectations about her. Despite the fact that game theory has been rendered mathematically and logically systematic only since 1944, game-theoretic insights can be found among commentators going back to ancient times. For example, in two of Plato's texts, the Laches and the Symposium, Socrates recalls an episode from the Battle of Delium that some commentators have interpreted (probably anachronistically) as involving the following situation. Consider a soldier at the front, waiting with his comrades to repulse an enemy attack. It may occur to him that if the defense is likely to be successful, then it isn't very probable that his own personal contribution will be essential. But if he stays, he runs the risk of being killed or wounded—apparently for no point. On the other hand, if the enemy is going to win the battle, then his chances of death or injury are higher still, and now quite clearly to no point, since the line will be overwhelmed anyway. Based on this reasoning, it would appear that the soldier is better off running away regardless of who is going to win the battle. But if all of the soldiers reason this way—as they all apparently should, since they're all in identical situations—then this will certainly bring about the outcome in which the battle is lost. Of course, this point, since it has occurred to us as analysts, can occur to the soldiers too. Does this give them a reason for staying at their posts? Just the contrary: the greater the soldiers' fear that the battle will be lost, the greater their incentive to get themselves out of harm's way. And the greater the soldiers' belief that the battle will be won, without the need of any particular individual's contributions, the less reason they have to stay and fight. If each soldier anticipates this sort of reasoning on the part of the others, all will quickly reason themselves into a panic, and their horrified commander will have a rout on his hands before the enemy has even engaged. Long before game theory had come along to show analysts how to think about this sort of problem systematically, it had occurred to some actual military leaders and influenced their strategies. Thus the Spanish conqueror Cortez, when landing in Mexico with a small force who had good reason to fear their capacity to repel attack from the far more numerous Aztecs, removed the risk that his troops might think their way into a retreat by burning the ships on which they had landed. With retreat having thus been rendered physically impossible, the Spanish soldiers had no better course of action than to stand and fight—and, furthermore, to fight with as much determination as they could muster. Better still, from Cortez's point of view, his action had a discouraging effect on the motivation of the Aztecs. He took care to burn his ships very visibly, so that the Aztecs would be sure to see what he had done. They then reasoned as follows: Any commander who could be so confident as to willfully destroy his own option to be prudent if the battle went badly for him must have good reasons for such extreme optimism. It cannot be wise to attack an opponent who has a good reason (whatever, exactly, it might be) for being sure that he can't lose. The Aztecs therefore retreated into the surrounding hills, and Cortez had the easiest possible victory. These two situations, at Delium and as manipulated by Cortez, have a common and interesting underlying logic. Notice that the soldiers are not motivated to retreat just, or even mainly, by their rational assessment of the dangers of battle and by their self-interest. Rather, they discover a sound reason to run away by realizing that what it makes sense for them to do depends on what it will make sense for others to do, and that all of the others can notice this too. Even a quite brave soldier may run rather than fight, but pointlessly die trying to stem the oncoming tide all by himself. Thus we could make the point, without contradiction, a circumstance in which an army, all of whose members are brave, flees at top speed before the enemy makes a move, if the soldiers really are brave, then this surely isn't the outcome any of them wanted; each would have preferred that all stand and fight. What we have here, then, is a case in which the interaction of many individually rational decision-making processes—one process per soldier—produces an outcome intended by no one. (Many armies try to avoid this problem just as Cortez did. Since they can't usually make retreat physically impossible, they make it economically irrational: for most of history, it was standard military practice to execute deserters. In that context standing and fighting is each soldier's individually rational course of action after all, because the expected cost of running is at least as high as the cost of staying.) Another classic source that invites this sequence of reasoning is found in Shakespeare's Henry V. During the Battle of Agincourt Henry decided to slaughter his French prisoners, in full view of the enemy and to the surprise of his subordinates, who describe the action as being out of moral character. The reasons Henry gives allude to non-strategic considerations: he is afraid that the prisoners may free themselves and threaten his position. However, a game theorist might have furnished him with supplementary strategic (and similarly prudential, though perhaps not moral) justification. His own troops observe that the prisoners have been killed, and observe that the enemy has observed this. Therefore, they know what fate will await them if the enemy's hand if they don't win. Metaphorically, but very effectively, their boats have been burnt. The slaughter of the prisoners plausibly sent a signal to the soldiers of both sides, thereby changing their incentives in ways that favoured English prospects for victory. These examples might seem to be relevant only for those who find themselves of cut-throat contention; perhaps, one might think, in medieval, stratagems, and military strategy. But the point is that the same logic, in different guises, is important in modern political philosophy, since centuries before anyone had an explicit name for this sort of logic. Philosophers share with social scientists the need to be able to represent and systematically model not only what they think people do, but what they often actually do in interactive situations. Hobbes's Leviathan is often regarded as the founding work in modern political philosophy, the text that began the continuing round of analyses of the function and justification of the state and its restrictions on individual liberties. The core of Hobbes's reasoning can be given straightforwardly as follows. The best situation for all people is one in which each is free to do as she pleases. (One may or may not agree with this as a matter of psychology or ideology, but it is Hobbes's assumption.) Often, such free people will wish to cooperate with one another in order to carry out projects that would be impossible for an individual acting alone. But if there are any immoral or amoral agents around, they will notice that their interests might at least sometimes be best served by getting the benefits from cooperation and not returning them. Suppose, for example, that you agree to help me build my house in return for my promise to help you build yours. After my house is finished, I can make your labour free to me simply by renegeing on my promise. I then realize, however, that if this leaves you with no house, you will have an incentive to take mine. This will put me in constant fear of you, and force me to spend valuable time and resources guarding myself against you. I can best minimize these costs by striking first and killing you at the first opportunity. Of course, you can anticipate all of this reasoning by me, and so have good reason to try to beat me to the punch. Since I can anticipate this reasoning by you, my original fear of you was not paranoid; nor was yours of me. In fact, neither of us actually needs to be immoral to get this chain of mutual reasoning going. The only thing that either of us needs to do is to expect the other to do the same. Once this is expected, neither of us needs to be immoral. The other might want, this murderous logic can take hold long before we are so silly as to imagine that we could ever actually get so far as making dead to help one another build houses in the first place. Left to their own devices, agents who are at least sometimes narrowly self-interested can repeatedly fail to derive the benefits of cooperation, and instead be trapped in a state of 'war of all against all', in Hobbes's words. In these circumstances, human life, as he vividly and famously put it, will be 'solitary, poor, nasty, brutish and short.' Hobbes's proposed solution to this problem was tyranny. The people can hire an agent—a government—whose job it is to punish anyone who breaks any promise. So long as the threatened punishment is sufficiently dire then the cost of renegeing on promises will exceed the cost of keeping them. The logic here is identical to that used by an army when it threatens to shoot deserters. If all people know that these incentives hold for most others, then cooperation will not only be possible, but can be the expected norm, so that the war of all against all becomes a general peace. Hobbes pushes the logic of this argument to a very strong conclusion, arguing that it implies not only a government with the right and the power to enforce cooperation, but an 'undivided' government in which the arbitrary will of a single ruler must impose absolute obligation on all. Few contemporary political theorists think that the particular steps by which Hobbes reasons his way to this conclusion are both sound and valid. Working through these issues here, however, would carry us away from our topic into details of contractarian political philosophy. What is important in the present context is that these details, as they are in fact pursued in contemporary debates, involve sophisticated interpretation of the issues using the resources of modern game theory (see, for example, Hampton 1986). Furthermore, Hobbes's most basic point, that the fundamental justification for the coercive authority and practices of governments is peoples' own need to protect themselves from other agents' well-being. Economists have idealized this set of assumptions purely as an idealization for purposes of analysis, not a possible state of affairs anyone could try (or should want to try) to institutionally establish. But until the mathematics of game theory matured near the end of the 1970s, economists had to hope that the more closely a model approximates perfect competition, the more efficient it will be. No such hope, however, can be mathematically or logically justified in general; indeed, as a strict generalization the assumption was shown to be false as far back as the 1950s. This article is not about the foundations of economics, but it is important for understanding the origins and scope of game theory to know that perfectly competitive markets have built into them a feature that renders them susceptible to parametric analysis. Because agents face no entry costs to markets, they will open shop in any given market until competition drives all profits to zero. This implies that if production costs are fixed and demand is exogenous, then agents have no options about how much to produce if they are trying to maximize the differences between their costs and their revenues. These production levels can be determined separately for each agent, so none need pay attention to what the others are doing; each agent treats her counterparts as passive features of the environment. The other kind of situation to which classical economic analysis can be applied without recourse to game theory is that of a monopoly facing many customers. Here, as long as no customer has a share of demand large enough to exert strategic leverage, non-parametric considerations drop out and the firm's task is only to identify the combination of price and production quantity at which it maximizes profit. However, both perfect and monopolistic competition are very special and unusual market arrangements. Prior to the advent of game theory, therefore, economists were severely limited in the class of circumstances to which they could straightforwardly apply their models. Economists should have been more professionally interested in the conditions and techniques for maximizing utility, rather than in the conditions and techniques for maximizing profit. The more generalization of welfare, in addition, played a special concern with economists, since it was the only way to avoid the problems of the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility' in a much more general way, and have used it increasingly to denote a more general, and increasingly clearer, notion of utility. The notion of utility was in the market proper, and the moral aspects of the market proper, and the moral aspects of the market proper are relevant to the most famous (but not the most typical) games, the so-called Prisoner's Dilemma, and to other, more typical games. In doing this, we will need to introduce, define and illustrate the basic elements and techniques of game theory. 2. Basic Elements and Techniques of Game Theory 2.1 Utility An economic agent is, by definition, an entity with preferences. Game theorists, like economists and philosophers who study practical choice, describe these by means of an abstract concept called utility. This refers to some ranking, on some specified scale, of the subjective welfare or change in subjective welfare that an agent derives from an event. By 'welfare' we refer to some normative ideal of relative alignment between states of the world and agents' valuations of the states in question, justified by reference to some background framework. For example, we might evaluate the relative welfare of countries (which we might model as agents for some purposes) by reference to their per capita incomes, and we might evaluate the relative welfare of an animal, in the context of predicting and explaining its behavioral dispositions, by reference to its expected evolutionary fitness. In the case of people, it is most typical in economics and applications of game theory to evaluate their relative welfare by reference to their own implicit or explicit judgments of it. This is why we referred above to subjective welfare. Consider a person who adores the taste of pickles but dislikes onions. She might be said to associate higher utility with states of the world in which, all else being equal, she consumes more pickles and fewer onions than with states in which she consumes more onions and fewer pickles. Examples of this kind suggest that 'utility' denotes a measure of subjective psychological fulfillment, and this is indeed how the concept was originally used by von Neumann and Morgenstern (1944). Economists have used the term 'utility'

[illegible]

one means they face no threat of punishment in the third-last round, and defect there too. We can simply iterate this backwards through the game tree until we reach the first round. Since cooperation is not a NE strategy in that round, tit-for-tat is no longer a NE strategy in the repeated game, and we get the same outcome—mutual defection—as in the one-shot PD. Therefore, cooperation is only possible in repeated PDs where the expected number of repetitions is indeterminate. (Of course, this does apply to many real-life games.) Note that in this context any amount of uncertainty in expectations, or possibility of trembling hands, will be conducive to cooperation, at least for awhile. When people in experiments play repeated PDs with known end-points, they indeed tend to cooperate for awhile, but learn to defect earlier as they gain experience. Now we introduce a third complication. Suppose that players' ability to distinguish defection from cooperation is imperfect. Consider our case of the widget cartel. Suppose the players observe a fall in the market price of widgets. Perhaps this is because a cartel member cheated. Or perhaps it has resulted from an exogenous drop in demand. If tit-for-tat players mistake the second case for the first, they will defect, thereby setting off a chain-reaction of mutual defections from which they can never recover, since every player will reply to the first encountered defection with defection, thereby begetting further defections, and so on. If players know that such miscommunication is possible, they have incentive to resort to more sophisticated strategies. In particular, they may be prepared to sometimes risk following defections with cooperation in order to test their inferences. However, if they are not prepared to do this, they are stuck in a chicken-and-egg problem: they need to coordinate in order to coordinate. In the case of the widget cartel, the cartel members need to coordinate on a mutually agreed-upon strategy for responding to a fall in the market price of widgets. This is a variable and complicated pattern of observable behavior), their use increases the probability of miscommunication. But miscommunication is what causes repeated-game cooperative equilibria to unravel in the first place. The complexities surrounding information signaling, screening and inference in repeated PDs help to intuitively explain the folk theorem, so called because no one is sure who first recognized it, that in repeated PDs, for any strategy (S) there exists a possible distribution of strategies among other players such that the vector of (S) and these other strategies is a NE. When critics of applications of game theory to behavioral and social science and business cases complain that the applications in question assume implausible levels of inferential capacity on the part of people, this is what they have in mind. In Section 5 we will consider a way of responding to this kind of concern. Real, complex, social and political dramas are seldom straightforward instantiations of simple games such as PDs. Hardin (1995) offers an analysis of two tragically real political cases, the Yugoslavian civil war of 1991–95, and the 1994 Rwandan genocide, as PDs that were nested inside coordination games. A coordination game occurs whenever the utility of two or more players is maximized by their doing the same thing as one another, and where such correspondence is more important to them than whatever it is, in particular, that they both do. A standard example arises with rules of the road: 'All drive on the left' and 'All drive on the right' are both outcomes that are NEs, and neither is more efficient than the other. In games of 'pure' coordination, it doesn't even help to use more selective equilibrium criteria. For example, suppose that we require our players to reason in accordance with Bayes's rule (see Section 3 above). In these circumstances, any strategy that is a best reply to any vector of mixed strategies available in NE is said to be rationalizable. That is, a player can find a set of systems of beliefs for the other players such that any history of the game along an equilibrium path is rationalizable. For example, in a traffic light coordination game, the NE is to proceed on green lights, but this is only rationalizable if the players believe that the other players will also proceed on green lights. If the players believe that they will proceed on red lights, then the NE is to proceed on red lights. This is a variable and complicated pattern of observable behavior), their use increases the probability of miscommunication. But miscommunication is what causes repeated-game cooperative equilibria to unravel in the first place. The complexities surrounding information signaling, screening and inference in repeated PDs help to intuitively explain the folk theorem, so called because no one is sure who first recognized it, that in repeated PDs, for any strategy (S) there exists a possible distribution of strategies among other players such that the vector of (S) and these other strategies is a NE. When critics of applications of game theory to behavioral and social science and business cases complain that the applications in question assume implausible levels of inferential capacity on the part of people, this is what they have in mind. In Section 5 we will consider a way of responding to this kind of concern. Real, complex, social and political dramas are seldom straightforward instantiations of simple games such as PDs. Hardin (1995) offers an analysis of two tragically real political cases, the Yugoslavian civil war of 1991–95, and the 1994 Rwandan genocide, as PDs that were nested inside coordination games. A coordination game occurs whenever the utility of two or more players is maximized by their doing the same thing as one another, and where such correspondence is more important to them than whatever it is, in particular, that they both do. A standard example arises with rules of the road: 'All drive on the left' and 'All drive on the right' are both outcomes that are NEs, and neither is more efficient than the other. In games of 'pure' coordination, it doesn't even help to use more selective equilibrium criteria. For example, suppose that we require our players to reason in accordance with Bayes's rule (see Section 3 above). In these circumstances, any strategy that is a best reply to any vector of mixed strategies available in NE is said to be rationalizable. That is, a player can find a set of systems of beliefs for the other players such that any history of the game along an equilibrium path is rationalizable. For example, in a traffic light coordination game, the NE is to proceed on green lights, but this is only rationalizable if the players believe that the other players will also proceed on green lights. If the players believe that they will proceed on red lights, then the NE is to proceed on red lights. This is a variable and complicated pattern of observable behavior), their use increases the probability of miscommunication. But miscommunication is what causes repeated-game cooperative equilibria to unravel in the first place. The complexities surrounding information signaling, screening and inference in repeated PDs help to intuitively explain the folk theorem, so called because no one is sure who first recognized it, that in repeated PDs, for any strategy (S) there exists a possible distribution of strategies among other players such that the vector of (S) and these other strategies is a NE. When critics of applications of game theory to behavioral and social science and business cases complain that the applications in question assume implausible levels of inferential capacity on the part of people, this is what they have in mind. In Section 5 we will consider a way of responding to this kind of concern. Real, complex, social and political dramas are seldom straightforward instantiations of simple games such as PDs. Hardin (1995) offers an analysis of two tragically real political cases, the Yugoslavian civil war of 1991–95, and the 1994 Rwandan genocide, as PDs that were nested inside coordination games. A coordination game occurs whenever the utility of two or more players is maximized by their doing the same thing as one another, and where such correspondence is more important to them than whatever it is, in particular, that they both do. A standard example arises with rules of the road: 'All drive on the left' and 'All drive on the right' are both outcomes that are NEs, and neither is more efficient than the other. In games of 'pure' coordination, it doesn't even help to use more selective equilibrium criteria. For example, suppose that we require our players to reason in accordance with Bayes's rule (see Section 3 above). In these circumstances, any strategy that is a best reply to any vector of mixed strategies available in NE is said to be rationalizable. That is, a player can find a set of systems of beliefs for the other players such that any history of the game along an equilibrium path is rationalizable. For example, in a traffic light coordination game, the NE is to proceed on green lights, but this is only rationalizable if the players believe that the other players will also proceed on green lights. If the players believe that they will proceed on red lights, then the NE is to proceed on red lights. This is a variable and complicated pattern of observable behavior), their use increases the probability of miscommunication. But miscommunication is what causes repeated-game cooperative equilibria to unravel in the first place. The complexities surrounding information signaling, screening and inference in repeated PDs help to intuitively explain the folk theorem, so called because no one is sure who first recognized it, that in repeated PDs, for any strategy (S) there exists a possible distribution of strategies among other players such that the vector of (S) and these other strategies is a NE. When critics of applications of game theory to behavioral and social science and business cases complain that the applications in question assume implausible levels of inferential capacity on the part of people, this is what they have in mind. In Section 5 we will consider a way of responding to this kind of concern. Real, complex, social and political dramas are seldom straightforward instantiations of simple games such as PDs. Hardin (1995) offers an analysis of two tragically real political cases, the Yugoslavian civil war of 1991–95, and the 1994 Rwandan genocide, as PDs that were nested inside coordination games. A coordination game occurs whenever the utility of two or more players is maximized by their doing the same thing as one another, and where such correspondence is more important to them than whatever it is, in particular, that they both do. A standard example arises with rules of the road: 'All drive on the left' and 'All drive on the right' are both outcomes that are NEs, and neither is more efficient than the other. In games of 'pure' coordination, it doesn't even help to use more selective equilibrium criteria. For example, suppose that we require our players to reason in accordance with Bayes's rule (see Section 3 above). In these circumstances, any strategy that is a best reply to any vector of mixed strategies available in NE is said to be rationalizable. That is, a player can find a set of systems of beliefs for the other players such that any history of the game along an equilibrium path is rationalizable. For example, in a traffic light coordination game, the NE is to proceed on green lights, but this is only rationalizable if the players believe that the other players will also proceed on green lights. If the players believe that they will proceed on red lights, then the NE is to proceed on red lights. This is a variable and complicated pattern of observable behavior), their use increases the probability of miscommunication. But miscommunication is what causes repeated-game cooperative equilibria to unravel in the first place. The complexities surrounding information signaling, screening and inference in repeated PDs help to intuitively explain the folk theorem, so called because no one is sure who first recognized it, that in repeated PDs, for any strategy (S) there exists a possible distribution of strategies among other players such that the vector of (S) and these other strategies is a NE. When critics of applications of game theory to behavioral and social science and business cases complain that the applications in question assume implausible levels of inferential capacity on the part of people, this is what they have in mind. In Section 5 we will consider a way of responding to this kind of concern. Real, complex, social and political dramas are seldom straightforward instantiations of simple games such as PDs. Hardin (1995) offers an analysis of two tragically real political cases, the Yugoslavian civil war of 1991–95, and the 1994 Rwandan genocide, as PDs that were nested inside coordination games. A coordination game occurs whenever the utility of two or more players is maximized by their doing the same thing as one another, and where such correspondence is more important to them than whatever it is, in particular, that they both do. A standard example arises with rules of the road: 'All drive on the left' and 'All drive on the right' are both outcomes that are NEs, and neither is more efficient than the other. In games of 'pure' coordination, it doesn't even help to use more selective equilibrium criteria. For example, suppose that we require our players to reason in accordance with Bayes's rule (see Section 3 above). In these circumstances, any strategy that is a best reply to any vector of mixed strategies available in NE is said to be rationalizable. That is, a player can find a set of systems of beliefs for the other players such that any history of the game along an equilibrium path is rationalizable. For example, in a traffic light coordination game, the NE is to proceed on green lights, but this is only rationalizable if the players believe that the other players will also proceed on green lights. If the players believe that they will proceed on red lights, then the NE is to proceed on red lights. This is a variable and complicated pattern of observable behavior), their use increases the probability of miscommunication. But miscommunication is what causes repeated-game cooperative equilibria to unravel in the first place. The complexities surrounding information signaling, screening and inference in repeated PDs help to intuitively explain the folk theorem, so called because no one is sure who first recognized it, that in repeated PDs, for any strategy (S) there exists a possible distribution of strategies among other players such that the vector of (S) and these other strategies is a NE. When critics of applications of game theory to behavioral and social science and business cases complain that the applications in question assume implausible levels of inferential capacity on the part of people, this is what they have in mind. In Section 5 we will consider a way of responding to this kind of concern. Real, complex, social and political dramas are seldom straightforward instantiations of simple games such as PDs. Hardin (1995) offers an analysis of two tragically real political cases, the Yugoslavian civil war of 1991–95, and the 1994 Rwandan genocide, as PDs that were nested inside coordination games. A coordination game occurs whenever the utility of two or more players is maximized by their doing the same thing as one another, and where such correspondence is more important to them than whatever it is, in particular, that they both do. A standard example arises with rules of the road: 'All drive on the left' and 'All drive on the right' are both outcomes that are NEs, and neither is more efficient than the other. In games of 'pure' coordination, it doesn't even help to use more selective equilibrium criteria. For example, suppose that we require our players to reason in accordance with Bayes's rule (see Section 3 above). In these circumstances, any strategy that is a best reply to any vector of mixed strategies available in NE is said to be rationalizable. That is, a player can find a set of systems of beliefs for the other players such that any history of the game along an equilibrium path is rationalizable. For example, in a traffic light coordination game, the NE is to proceed on green lights, but this is only rationalizable if the players believe that the other players will also proceed on green lights. If the players believe that they will proceed on red lights, then the NE is to proceed on red lights. This is a variable and complicated pattern of observable behavior), their use increases the probability of miscommunication. But miscommunication is what causes repeated-game cooperative equilibria to unravel in the first place. The complexities surrounding information signaling, screening and inference in repeated PDs help to intuitively explain the folk theorem, so called because no one is sure who first recognized it, that in repeated PDs, for any strategy (S) there exists a possible distribution of strategies among other players such that the vector of (S) and these other strategies is a NE. When critics of applications of game theory to behavioral and social science and business cases complain that the applications in question assume implausible levels of inferential capacity on the part of people, this is what they have in mind. In Section 5 we will consider a way of responding to this kind of concern. Real, complex, social and political dramas are seldom straightforward instantiations of simple games such as PDs. Hardin (1995) offers an analysis of two tragically real political cases, the Yugoslavian civil war of 1991–95, and the 1994 Rwandan genocide, as PDs that were nested inside coordination games. A coordination game occurs whenever the utility of two or more players is maximized by their doing the same thing as one another, and where such correspondence is more important to them than whatever it is, in particular, that they both do. A standard example arises with rules of the road: 'All drive on the left' and 'All drive on the right' are both outcomes that are NEs, and neither is more efficient than the other. In games of 'pure' coordination, it doesn't even help to use more selective equilibrium criteria. For example, suppose that we require our players to reason in accordance with Bayes's rule (see Section 3 above). In these circumstances, any strategy that is a best reply to any vector of mixed strategies available in NE is said to be rationalizable. That is, a player can find a set of systems of beliefs for the other players such that any history of the game along an equilibrium path is rationalizable. For example, in a traffic light coordination game, the NE is to proceed on green lights, but this is only rationalizable if the players believe that the other players will also proceed on green lights. If the players believe that they will proceed on red lights, then the NE is to proceed on red lights. This is a variable and complicated pattern of observable behavior), their use increases the probability of miscommunication. But miscommunication is what causes repeated-game cooperative equilibria to unravel in the first place. The complexities surrounding information signaling, screening and inference in repeated PDs help to intuitively explain the folk theorem, so called because no one is sure who first recognized it, that in repeated PDs, for any strategy (S) there exists a possible distribution of strategies among other players such that the vector of (S) and these other strategies is a NE. When critics of applications of game theory to behavioral and social science and business cases complain that the applications in question assume implausible levels of inferential capacity on the part of people, this is what they have in mind. In Section 5 we will consider a way of responding to this kind of concern. Real, complex, social and political dramas are seldom straightforward

[illegible]

as a way of discouraging examination of more egalitarian bargaining equilibria that are within reach along societies' equilibrium paths. Resolution of this debate between Gintis and Binmore significantly need not wait upon discoveries about the deep human evolutionary past that we may never have. The models make rival empirical predictions of some testable phenomena. If Gintis is right then there are limits, imposed by the discontinuity in hominin evolution, on the extent to which people can learn to be self-regarding. This is the main significance of the controversy discussed above over Heinrich et al.'s interpretation of their field data. Binmore's model of social equilibrium selection also depends, unlike Gintis's, on widespread dispositions among people to inflict second-order punishment on members of society who fail to sanction violators of social norms. Gintis (2005) shows using a game theory model that this is implausible if punishment costs are significant. However, Ross (2008a) argues that the widespread assumption in the literature that punishment of norm-violation must be costly results from failure to adequately distinguish between models of the original evolution of sociality, on the one hand, and models of the maintenance and development of norms and institutions once an initial set of them has stabilized. Finally, Ross also points out that Binmore's objectives are as much normative as descriptive: he aims to show egalitarians how to diagnose the errors in conservative rationalisations of the status quo without calling for revolutions that put equilibrium path stability (and, therefore, social welfare) at risk. It is a sound principle in constructing reform proposals that they should be 'knave-proof' (as Hume put it), that is, should be compatible with less altruism than might prevail in people.

9. Looking Ahead: Areas of Current Innovation

In 2016 the Journal of Economic Perspectives published a symposium on "What is Happening in Game Theory?" Each of the participants noted independently that game theory has become so tightly entangled with microeconomic theory in general that the question becomes difficult to distinguish from inquiry into the moving frontier of that entire sub-discipline, which is in turn the largest part of economics as a whole. Thus the boundary between the philosophy of game theory and the philosophy of microeconomics is now similarly indistinct. Of course, as has been stressed, applications of game theory extend beyond the traditional domain of economics, into all of the behavioral and social sciences. But as the methods of game theory have fused with the methods of microeconomics, a commentator might equally view these extensions as being exported applications of microeconomics. Following decades of development (incompletely) surveyed in the present article, the past few years have been relatively quiet ones where foundational innovations of the kind that invite contributions from philosophers are concerned. Some parts of the original foundations are being newly revisited, however: von Neumann and Morgenstern's (1944) introduction of game theory divided the inquiry into two parts. Noncooperative game theory analyzes cases built on the assumption that each player maximizes her own utility function while treating the expected strategic responses of other players as constraints. As discussed above, the specific game to which von Neumann and Morgenstern applied their modeling was poker, which is a zero-sum game. Most of the present article has focused on the many theoretical challenges and insights that arose from extending noncooperative game theory beyond the zero-sum domain. But this in fact develops only half of von Neumann and Morgenstern's classic. The other half developed cooperative game theory, about which nothing has so far been said here. The reason for this silence is that for most game theorists cooperative game theory is a distraction at best and at worst a technology that confuses the point of game theory by bypassing the aspect of games that mainly makes them potentially interesting and insightful in application, namely, the requirement that equilibria be selected endogenously under the restrictions imposed by Nash (1950a). This, after all, is what makes equilibria self-enforcing, just in the way that prices in competitive markets are, and thus renders them stable unless shocked from outside. Nash (1953) argued that solutions to cooperative games should always be verified by showing that they are also solutions to formally equivalent noncooperative games. Nash's accomplishment in the paper was the analytical identification of the relevant equivalence. One way of interpreting this was as demonstrating the ultimate redundancy of cooperative game theory. Cooperative game theory begins from the assumption that players have already, by some unspecified process, agreed on a vector of strategies, and thus on an outcome. Then the analyst deploys the theory to determine the minimal set of conditions under which the agreement remains stable. The idea is typically illustrated by the example of a parliamentary coalition. Suppose that there is one dominant party that must be a member of any coalition if it is to commend a majority of parliamentary votes on legislation and confidence. There might then be a range of alternative possible groupings of other parties that could sustain it. Imagine, to make the example more structured and interesting, that some parties will not serve in a coalition that includes certain specific others; so the problem faced by the coalition organizers is not simply a matter of summing potential votes. The cooperative game theorist identifies the set of possible coalitions. There may be some other parties, in addition to the dominant party, that turn out to be needed in every possible coalition. Identifying these parties would, in this example, reveal the core of the game, the elements shared by all equilibria. The core is the key solution concept of cooperative game theory, for which Shapley shared the Nobel prize. (Shapley (1953) is the great paper.) Nash (1953) defined the "Nash program" as consisting of verifying a particular cooperative equilibrium by showing that noncooperative players could arrive at it through the sequential bargaining process specified in Nash (1950b), and that all outcomes of such bargaining would include the core. In light of the example, it is no surprise that political scientists were the primary users of cooperative theory during the years while noncooperative game theory was still being fully developed. It has also been applied usefully by labor economists studying settlement negotiations between firms and unions, and by analysts of international trade negotiations. We might illustrate the value of such application by reference to the second example. Suppose that, given the weight of domestic lobbies in South Africa, the South African government will never agree to any trade agreement that does not allow it to protect its automotive assembly sector. (This has in fact been the case so far.) Then allowance for such protection is part of the core of any trade treaty another country or bloc might conclude with South Africa. Knowing this can help the parties during negotiations avoid rhetoric or commitments to other lobbies, in any of the negotiating countries, that would put the core out of reach and thus guarantee negotiation failure. This example also helps us illustrate the limitations of cooperative game theory. South Africa will have to trade off the interests of some other lobbies to protect its automotive industry. Which others will get traded off will be a function of the extensive-form play of non-cooperative sequential proposals and counter-proposals, and the South African bargainers, if they have done their due diligence, must be attentive to which paths through the tree throw which specific domestic interests under the proverbial bus. Thus carrying out the cooperative analysis does not relieve them of the need to also conduct the noncooperative analysis. Their game theory consultants might as well simply code the non-cooperative parameters into their Gambit software, which will output the core if asked. But cooperative game theory did not die, or become confined to political science applications. There has turned out to be a range of policy problems, involving many players whose attributes vary but whose ordinal utility functions are symmetrical, for which noncooperative modeling, while possible in principle, is absurdly cumbersome and computationally demanding, but for which cooperative modeling is beautifully suited. That we be dealing with ordinal utility functions is important, because in the relevant markets there are often no prices. The classic example (Gale and Shapley 1962) is a marriage market. Abstracting from the scale of individual romantic dramas and comedies, society features, as it were, a vast set of people who want to form into pairs, but care very much who they end up paired with. Suppose we have a finite set of such people. Imagine that the match-maker, or app, first splits the set into two proper subsets, and announces a rule that everyone in subset \A) will propose to someone in subset \B). Each of those in \B) who receive a proposal knows that she is the first choice of someone in \A). She selects her first choice from the proposals she has received and throws the rest back into the pool. Those in \A) whose initial proposals were not accepted now each propose to someone they did not propose to before, but possibly including people who are holding proposals from a previous round—Nkosi knows that Barbara preferred Amalia in round 1, but Nkosi wasn't part of that choice set and so might displace Amalia in round 2). Provably there exists a terminal round after which no further proposals will be made, and the matchmaking app will have found the core of the cooperative game because no person \i) in set \B) will prefer to pair with someone from set \A) who prefers \i) to whoever is holding that \A)-set dreamboat's proposal. Everyone from set B will now accept the proposal they are holding, and, if the two sets had the same cardinality and everyone would rather pair with someone than pair with no one, then nobody will go off alone. This is not a directly applicable model of a marriage market, so there is no money to be made in selling the simple matchmaking app described above. The problem is that we have no guarantee that, in the example, Nkosi and Amalia aren't one another's partners of destiny, but cannot get paired because they both began in subset \A). In game theory textbooks this problem is often finessed by assuming that Set \A) contains men and Set \B) contains women, and that everyone is so committed to heterosexuality that they'd rather pair with anyone of the opposite sex than anyone of their own sex. On the other hand, the model provides some insight, in the way that models typically do, if we don't insist on applying it too literally. After working through it, one sees the logic of facts about society that someone designing a real matchmaking app had better understand: that the app will have to log proposals under consideration but not yet accepted, leave people holding proposals under consideration on the market, and remember who has previously rejected whom (without creating a generalised emotional catastrophe by publicly posting this information). The real app will not be able to reliably find the core of the cooperative game, unless the set of people in the market is small, restricted, and has self-sorted into subsets to at least some extent by providing such information as "\X)-type person seeks \Y)-type person" for \X) and \Y) properties that everyone prioritizes. (Are there such properties, at least as an approximation?) But the real matchmaking apps seem to work well enough to be transforming the way in which most young people now find mates in countries with generally available internet access. Relationships between theoretically idealized and real marriage markets are comprehensively reviewed in Chiappori (2017). The revival of cooperative game theory as site of renewed interest has occurred because policy problems have been encountered that, unlike the original toy illustration using the all-straights marriage market, satisfy the model's crucial assumptions. Leading instances are matching university applicants and universities, and matching people needing organ transplants with donors (see Roth 2015). In these markets, there is no ambivalence about partitioning the sets to be matched. Ordinal preferences are the relevant ones: universities don't auction off places to the highest bidder (or at least not in general), and organs are not for sale (or at least not legally). The models are really applied, and they demonstrably have improved efficiency and saved lives. It is common in science for models that are practically clumsy fits to their original problems to turn out to furnish highly efficient solutions to new problems thrown up by technological change. The internet has created an environment for applications of matching algorithms—travellers and flat renters, diners and restaurants, students and tutors, and (regrettably) socially alienated people and purveyors of propaganda and fanaticism—that could have been designed by a theorist at any time since Shapley's original innovations, but would previously have been practically impossible to implement. These applications of cooperative game theory are often applied conjointly with the noncooperative game theory of auctions (Klemperer 2004) to drive market designs for goods and services so efficient as to be annihilating the once mighty shopping mall in even the suburban USA. Why are hotels more profitable and easily available than was the case in all but the largest cities before about 2007? The answer is that dynamic pricing algorithms (Gershkov and Moldovanu 2014) blend matching theory and auction theory to allow hotels, combined with online travel service aggregators, to find customers willing to pay premium rates for their ideal locations and times, and then fill the remaining rooms with bargain hunters whose preferences are more flexible. Airlines operate similar technology. Game theory thus continues to be one of the 20th-century inventions that is driving social revolutions in the 21st, and Samuelson (2016) predicts a coming surge of renewed interest in the deeper mathematics of cooperative games and their relationships to noncooperative games. A range of further applications of both classical and evolutionary game theory have been developed, but we have hopefully now provided enough to convince the reader of the tremendous, and constantly expanding, utility of this analytical tool. The reader whose appetite for more has been aroused should find that she now has sufficient grasp of fundamentals to be able to work through the large literature, of which some highlights are listed below.