Appendix A A Primer in Game Theory

This presentation of the main ideas and concepts of game theory required to understand the discussion in this book is intended for readers without previous exposure to game theory. [1]

A game-theoretic analysis starts by specifying the rules of the game. These rules identify the decision makers (the players), their possible actions, the information available to them, the probability distributions over chance events, and each decision maker's preference over outcomes, specifically, the set of all possible combinations of actions by the players. A game is represented, or defined, by the triplet of the players' set, the action set (which specifies each player's actions), and the payoff set (which specifies each player's payoffs as a function of the actions taken by the players). The rules of the game are assumed to be common knowledge.[2] The strategy set in a game is the set of all possible plans of actions by all the players when each conditions her action on the information available to her. The situations considered are strategic in the sense that each player's optimal strategy depends on the actions of other players. (Nonstrategic situations constitute a special case.)

The objective of game-theoretic analysis is to predict behavior in strategic situations—to predict an action combination (an action to each player) for any given rules of the game. The difficulty of finding such solutions stems from the fact that because the action optimal for each player depends on others' actions, no player can choose his optimal action independently of what other players do. For player A to choose behavior, he has to know what B will do, but for B to choose behavior, he has to know what A will do. The classical game-theoretic concepts of Nash equilibrium and its refinements, such as subgame perfect equilibrium, mitigate this infinite loop problem and eliminate some action combinations as implausible in a given game.

---

[1] For a relatively nontechnical introduction to game theory, see Dixit and Nalebuff (1991); Gibbons (1992, 1998); and Watson (2001). For a more technical analysis, see Fudenberg and Tirole (1991) and Gintis (2000). See Aumann and Hart (1994, 2002) for an extensive review of the application of game theory to economics and political science; Milgrom and Roberts (1995) on organizational theory; Hart and Holmstrom (1987) and Hart (1995) on contract theory; and Weingast (1996), Sened (1997), and Bates et al. (1998) on political science.

[2] S is common knowledge if all players know S, all players know that all players know S, and so on ad infinitum (D.Lewis 1969). In games of complete information, the rules of the game are common knowledge. In games of incomplete information, the probability distribution of the aspect of the game that is not common knowledge is common knowledge.

The basic idea of the Nash restriction is not to consider the dynamic problem of choosing behavior but to consider behavior that constitutes a solution to the problem of choosing behavior. Nash equilibria restrict admissible solutions (action combinations) to those that are self-enforcing: if each individual expects others to follow the behavior expected of them, he finds it optimal to follow the behavior expected of him.

To keep the discussion simple, I concentrate, without loss of generality, on two-player games, although the analysis applies to games with more players as well. Sections A.1 and A.2 examine static games in which the players move simultaneously and dynamic games in which the players move sequentially, respectively. Section A.3 then discusses repeated game theory, which examines situations in which a particular stage game, either static or dynamic, is repeated over time. Knowledge of games with incomplete information, in which players have different information regarding aspects of the structure of the game, is not essential for reading the book. Short discussions of such games are provided in Chapter 3 and Appendix C, section C.1. Chapter 5 discusses learning game theory, while Appendix C section C2.7, discusses imperfect monitoring.

A.1 Self-Enforcing Behavior in Static Games: The Nash Equilibrium

Consider first static (or simultaneous-move) games—games in which all players take actions simultaneously. Assume that all players have the same information about the situation. The structure of such games is as follows: Player 1 chooses an action a_1 from the set of feasible actions A_1. Simultaneously, player 2 chooses an action a_2 from the set of feasible actions A_2. After the players choose their actions, they receive the following payoffs: $u_1(a_1, a_2)$ to player 1 and $u_2(a_1, a_2)$ to player 2.

The prisoners' dilemma game is perhaps the best-known and most explored static game. It is so well known because it illustrates that in strategic situations, rationality alone is insufficient to reach a Pareto-optimal outcome. Unlike in market situations, in strategic situations one's desire to improve his lot does not necessarily maximize social welfare. In the prisoners' dilemma game, each player can either cooperate with the other or defect. If both cooperate, each player's payoff will be higher than if they both defect. But if one cooperates and the other defects, the defector benefits, receiving a higher payoff than if both cooperate. Meanwhile, the cooperator receives a lower payoff than he would have had he also defected.

Figure A.1 presents a particular prisoners' dilemma game. The players' actions are denoted by C (cooperate) and D (defect). Each cell corresponds to an action combination, or a pair of actions. The payoffs associated with each action combination are represented by two numbers, the payoff to player 1 and the payoff to player 2.

	Player 2's actions		
Player 1's actions		C	D
	C	1, 1	-15, 5
	D	5, -15	-8, -8

Figure A.1. The Prisoners' Dilemma Game

In this game, the best each player can do is defect. Player 1 cannot expect player 2 to play C, because no matter what player 1 does, player 2 is better off playing D. If player 1 plays C, then player 2 gains 1 playing C but 5 playing D. If player 1 plays D, then player 2 gains -15 from playing C and only -8 from playing C. The same holds for player 1, who is always better off playing D. In the language of game theory, defecting is each player's dominant strategy: it is the best that he can do, independent of what the other player does. Hence the action combination (D, D) will be followed if the game captures all aspects of the situation.

In the particular case of the prisoners' dilemma, one's expectations about the behavior of the other player do not matter when choosing an action. Playing D is the best one can do regardless of the other's choice of action. But in strategic situations in general, a player's optimal choice of action depends on the other player's choice of action.

Consider, for example, the driving game presented in Figure A.2. This game represents a situation in which two drivers are heading toward each other. Both players can choose to drive on the left or the right side of the road. If they both choose the same side, a collision is avoided and each receives a payoff of 2. If they choose opposite sides, either (right, left) or (left, right), they collide, and each receives a payoff of 0.

	Player 2's actions		
Player 1's actions		Left	Right
	Left	2, 2	0, 0
	Right	0, 0	2, 2

Figure A.2. The Driving Game

In this game the situation is strategic: the best action for one player depends on the action of the other. If player 1 is expected to choose left, for example, player 2's optimal response is to play left, thereby earning 2 instead of 0 from playing right. But if player 1 is expected to play right, player 2 is better off playing right as well. Player 2's optimal choice depends on player 1's actual choice. To choose an action, player 2 has to know the action of player 1. But the same holds for player 1. As each player's choice of action depends on that of that of the other, neither can choose an action.

This interrelatedness of decisions implies that we cannot find out what the players will do by examining the behavior of each of them separately, as we did in the prisoners' dilemma game. The ingenuity of the Nash equilibrium concept is that instead of attempting to find out what the players will do by examining the players' decision processes, we find possible outcomes by considering what outcomes if expected will be followed.

Suppose that it is common knowledge that both players hold the same expectations about how the game will be played. What expectation about behavior can they hold? They can expect only that self-enforcing behavior will be followed. Behavior is self-enforcing if, when players expect it to be followed, it is indeed followed because each player finds it optimal to do so expecting the others to follow it. An action combination (often referred to also as a strategy combination) satisfying this condition is called a Nash equilibrium. A Nash equilibrium fulfills a mutual best-response condition: each player's best response to his correct beliefs regarding the others' behavior is to follow the behavior expected of him.[3]

[3] In static games an action combination (a_1^*, a_2^*) is a Nash equilibrium if a_1^* is a best response for player 1 to a_2^* and a_2^* is a best response to a_1^*. That is, a_1^* must satisfy $u_1(a_1^*, a_2^*) \geq u_1(a_1, a_2^*)$ for every a_1 in A_1, and $u_2(a_1^*, a_2^*) \geq u_2(a_1^*, a_2)$ for every a_2 in A_2.

To illustrate that not all behavior satisfies this condition, consider behavior that, if expected, will not be followed. In the driving game, this case occurs with respect to the action combination (right, left). This combination would not be followed if each player expected the other to follow it. If player 2 expects player 1 to play right, her best response is to play right, receiving 2 instead of 0. Hence player 1 cannot hold the belief that player 2 will play left in this case. We can continue to consider whether various action combinations are self-enforcing in this manner. This analysis yields that the driving game has two Nash equilibria, (left, left) and (right, right).[4] If, for example, (left, left) is expected, both players will find it optimal to drive on the left because, expecting the other to do so, it is each driver's best response. Indeed, each of these Nash equilibria prevails in different countries. This analysis also illustrates that a game can have multiple Nash equilibria.

Some games do not have an action combination that satisfies the Nash condition. Consider the matching pennies game in Figure A.3. Each of the two players simultaneously chooses either head or tail. If their choices do not match, player 2 loses receiving -1 while player 1 receives 1. If they do match, player 1 loses receiving -1 while player 2 receives 1. In this game, there is no Nash equilibrium, as defined previously. This lack of an equilibrium reflects that this game captures a situation in which each player tries to outguess the action of the other. If player 1 expects player 2 to play heads (tails), his best response is to play tails (heads).

		Player 2's actions	
		Head	Tail
Player 1's actions	Head	-1, 1	1, -1
	Tail	1, -1	-1, 1

Figure A.3. The Matching Pennies Game

It is reasonable in such situations that peoples' expectations about behavior will be probabilistic in nature. People will expect others to play heads some of the time and tails some of the time. Game theory defines Nash equilibrium in such cases as well. This is done by referring

[4] There is also a third, mixed-strategy Nash equilibrium, in which each player chooses which side to drive on with probability 0.5. See the discussion of this notion better in this appendix.

to the actions in a player's action set (A_i) as pure strategies and defining a mixed strategy as a probability distribution over the player's pure strategies. We can then solve for the so-called mixed-strategy Nash equilibrium.[5] In the matching pennies game and the driving game, for example, playing each action with a probability of 0.5 for each player is a mixed-strategy Nash equilibrium.

Any game with a finite number of players, each of whom has a finite number of pure strategies, has a Nash equilibrium, although possibly only in mixed strategies. By restricting action combinations (i.e., plans of behavior) to those that are self-enforcing in the Nash equilibrium sense, game theory restricts the set of admissible behavior in such games.

Although the situations described here are very simple, the same analysis can be applied to more complicated ones, in which players move sequentially and there is asymmetric information or uncertainty. The equilibrium notions used for such situations are, by and large, refinements of the Nash equilibrium - that is, they are Nash equilibria that fulfill some additional conditions. The following discussion of dynamic games illustrates the nature of these refinements and the usefulness of imposing further restrictions on admissible self-enforcing behavior.

A.2 Self-Enforcing Behavior in Dynamic Games: Backward Induction and Subgame Perfect Equilibria

Consider a dynamic situation in which the players move sequentially rather than simultaneously. It is easier to present dynamic games in extensive (tree-diagram) form than in the normal (matrix) form used in Figures A.1-3. In extensive form a game is presented as a graph or a tree in which a branching point is a decision point for a player and each branch is associated with a different action. The payoffs associated with different actions are denoted at the end of the tree.

Although dynamic games can have many branches and decision points, their basic structure can be illustrated in the case of a game with two decision points. In this game player 1 chooses an action a_1 from the set of feasible actions A_1. After observing player 1's choice, player 2 chooses an action a_2 from the set of feasible actions A_2. After the players choose their actions, they receive payoffs $u_1(a_1, a_2)$ to player 1 and $u_2(a_1, a_2)$ to player 2.

[5] Harsanyi provided an interpretation of this mixing as reflecting one's uncertainty about the other player's choice of action. For an intuitive account, see Gibbons (1998).

The one-sided prisoner’s dilemma game is an example of a dynamic game with this structure (Figure A.4). First, player 1 chooses either to cooperate or defect. If he chooses to defect, the game ends and the players’ payoffs are (.5, 0). If player 1 chooses to cooperate, player 2 can choose an action. If he chooses to cooperate, both players’ payoffs are 1, but if he chooses to cheat, he receives the higher payoff of 2, while player 1 receives a payoff of 0.[6] In this game, player 1 can gain from cooperating, but only if player 2 cooperates. If player 2 cheats, player 1 receives a lower payoff than if he had not cooperated.

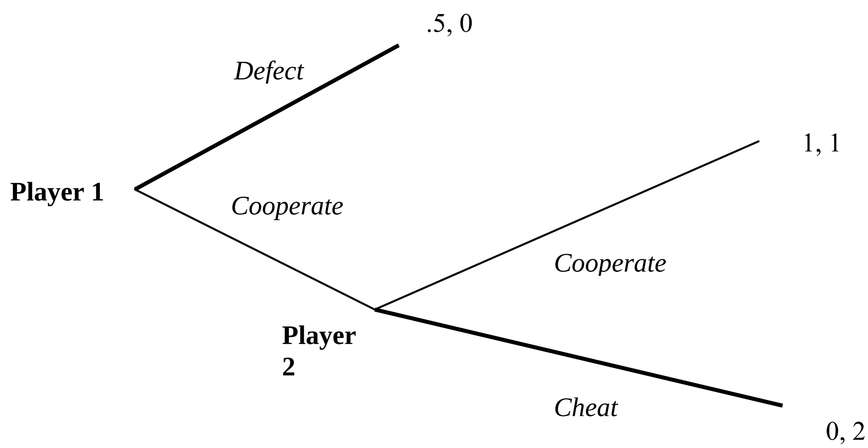

Figure A.4. The One-Sided Prisoner’s Dilemma Game

Dynamic games such as the one-sided prisoner’s dilemma are of interest in the social sciences because they capture an essential part of all exchange relationships—personal, social, economic, and political. Exchange is always sequential: some time elapses between the *quid* and the *quo* (Greif 1997a; 2000). More generally, in social relationships one often has to give before receiving; at the moment of giving, one receives only a promise of receiving something in the future.

[6] This game is also known as the game of trust (Kreps 1990a). Player 1 can either not trust (defect) or trust (cooperate). If player 1 does not trust, the game is over. If he trusts, player 2 can decide whether to honor the trust (cooperate) or to renege (cheat).

Can player 1 trust player 2 to cooperate? To find out, we can work backward through the game tree, examining the optimal action of the player who is supposed to move at each branching point.[7] This method is known as backward induction.

Consider player 2's decision. He receives a payoff of 2 from cheating and a payoff of 1 from cooperating, implying that cheating is his optimal choice. Expecting that, player 1 will choose to defect and receive .5 rather than cooperate and receive 0. (These branches are in bold in the game tree diagram in figure A.4.) This action combination is self-enforcing, because player 1's best response to reneging is to defect, while player 2's best response to defecting is to cheat. Backward induction reveals the self-enforcing action combination of (defect, cheat). This action combination is a Nash equilibrium.

As this analysis indicates, Nash equilibria can be Pareto-inferior. The payoffs associated with (cooperate, cooperate) leave each player better off than if player 1 defects; cooperation is thus profitable and efficient. But if player 1 cooperates, the payoff to player 2 from cheating is higher than from cooperating. Cooperation is not self-enforcing.

In the one-sided prisoner's dilemma game, backward induction yields the only Nash equilibrium. This can easily be seen if we present the game in matrix form (Figure A.5). In matrix form, player 1 chooses between cooperating and defecting, while player 2 chooses between cooperating and reneging. The payoffs associated with each action combination are the same as those in Figure A.4. The Nash equilibrium outcome is in boldface.

		Player 2's actions	
		Cooperate	Cheat
Player 1's actions	Cooperate	1, 1	0, 2
	Defect	.5, 0	**.5, 0**

Figure A.5. One-Sided Prisoner's Dilemma Game in Matrix Form

[7] For experimental evidence on people's use of backward induction, see Appendix B. For the theoretical weaknesses of backward induction and subgame perfection, see Fudenberg and Tirole (1991); Binmore (1996); and Hardin (1997).

When backward induction is possible, it always leads to action combinations that are Nash equilibria, but the opposite does not hold. If we represent an extensive (tree-diagram) form in the associated matrix (normal) form, not every Nash equilibrium in the game's matrix form can be reached through backward induction in the original tree form. This is because analyzing the game in tree form using backward induction captures that the players move sequentially, something that is not captured in the matrix form representation of the game. That the tree form captures more information about the structure of the game allows us to eliminate some Nash equilibria that we cannot eliminate in the normal form. Specifically, we can eliminate Nash equilibria that are based on noncredible threats or promises. The tree representation thus assists in deductively restricting—refining—the set of admissible self-enforcing behavior.

To see this advantage of backward induction, consider the following tree and matrix presentations of the same game (Figure A.6). In this game, player 1 chooses between playing left (L) or right (R), while player 2, who moves second, chooses between playing up (U) or down (D). If player 1 plays L, the payoffs are 1 to player 1 and 2 to player 2. If player 1 plays R and player 2 plays D, the payoffs are (2, 1) but if player 2 plays U, the payoffs are (0, 0.) The analysis of this game illustrates how backward induction eliminates Nash equilibria based on noncredible threats.

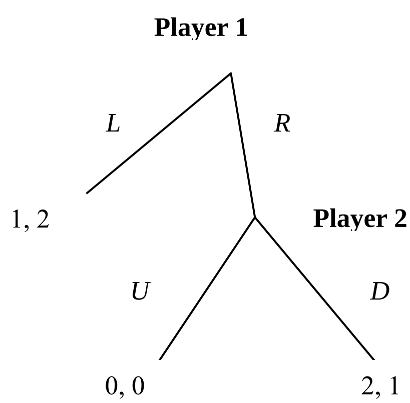

		Player 2's actions	
		U	*D*
Player 1's actions	*L*	1, 2	1 ,2
	R	0, 0	2, 1

Figure A.6. Elimination of Nash Equilibria Based on Noncredible Threats through Backward Induction

The matrix form presentation of this game shows two Nash equilibria: (L, U), with payoff (1, 2,) and (R, D), with payoff (2, 1). Backward induction yields only (R, D). (L, U) did not survive backward induction, because it relies on a noncredible threat that is concealed by the normal form presentation. In this equilibrium, player 1 is motivated to choose L because player 2 is supposed to play U, while player 2's best response to player 1's choice of L is indeed U. Given that player 1 chose L, player 2's payoff does not really depend on choosing between U and D, because given that player 1 chose L neither of these actions would be taken. Hence the equilibrium (L, U) depends on a noncredible threat off the equilibrium path—that is, it relies on player 2 taking an action in a situation that would never occur if the players play according to this action combination. Had the need for player 2 to take this action actually risen, he would not have found it optimal to do so. Backward induction enables us to call player 2's bluff and restrict the set of admissible self-enforcing behavior accordingly. If player 1 played R and hence player 2's choice of action influences the payoffs, playing D and receiving 1 (instead of playing U and receiving 0) is optimal for player 2. Backward induction captures that player 1, anticipating that response,would choose R and receive 2 rather than choose L and receive 1.

Backward induction can be applied in any dynamic finite-horizon game of perfect information. In such games the players move sequentially and all previous moves become common knowledge before the next action has to be chosen. In other games, such as dynamic games with simultaneous moves or an infinite horizon, however, we cannot apply backward induction directly. The notion of subgame perfect equilibrium enables us nevertheless to restrict the set of admissible Nash equilibrium by eliminating those that rely on noncredible threats or promises. Indeed, when backward induction can be applied, the resulting Nash equilibrium is a subgame perfect equilibrium—it is a refinement of Nash equilibrium in the sense that it is a Nash equilibrium that satisfies an additional requirement.

To grasp the concept of subgame perfect equilibrium intuitively, note that in the examples presented here, the action combinations yielded by backward induction satisfied the mutual-best-response requirement of Nash equilibrium. It also satisfied the requirement that player 2's action be optimal in the game that begins when he has to choose an action. Beginning at this decision point, backward induction restricts the admissible action of player 2 to be optimal.

In dynamic games with simultaneous moves, however, we cannot, in general, follow this procedure, because an optimal action depends on the action of the other player. To see why this condition limits the use of backward induction, consider the following game, presented in both extended and normal form (Figure A.7). Player 1 moves first, choosing between A and B. If player 1 chooses B, the game is over and the payoffs are (2, 6). If player 1 chooses A, both players play the simultaneous move game presented in the two-by-two matrix. In the two-by-two game that follows player 1's choice of action A, backward induction cannot be applied by considering the optimal moves of either player 1 or player 2. Each player's optimal action depends on the action of the other. In other words, no player moves last, as in a sequential move game.

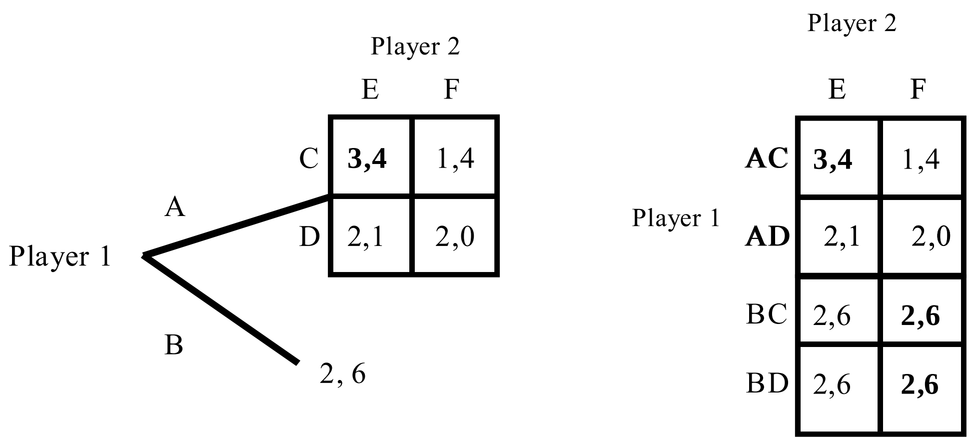

Figure A.7. Subgame Perfection

We can still, however, follow the logic of the backward induction procedure by finding the Nash equilibrium in the two-by-two game and considering player 1's optimal choice between A and B, taking this Nash equilibrium outcome into consideration. The Nash equilibrium in the two-by-two game is (C, E), which yields the payoffs (3, 4). Player 1's optimal choice between A and B is therefore A. The action combination that this procedure yields is (AC, E), which is a subgame perfect equilibrium.

To see that this procedure eliminates Nash equilibria that rely on noncredible threats, note that there are three Nash equilibria in the game: (AC, E), (BC, F), and (BD, F). (BC, F) and (BD, F) yield payoffs of (2, 6), making player 1 worse off and player 2 better off than the (AC, E) subgame perfect equilibrium. Both of these equilibria, however, rely on noncredible threats off the equilibrium path. Consider (BC, F). While considering the game as a whole, the choice of C or F does not affect payoffs, because these actions are off the path of play. But if the need to actually take these actions arises, they would not constitute a mutual best response. If player 2 chooses F, player 1's best response is D rather than C, which yields player 1 a payoff of 2 instead of 1. Similarly, in (BD, F), if player 1 chooses D, player 2's best response is E instead of F, which yields him a payoff of 1 instead of 0.

The notion of a subgame perfect equilibrium applies the mutual-best-response idea that is the essence of Nash equilibrium to subgames. Intuitively, a subgame is part of the original game that remains to be played, but a subgame begins only at points at which the complete history of how the game was already played is known to all players. A Nash equilibrium (in the game as a whole) is a subgame perfect equilibrium if the players' strategies constitute a Nash equilibrium in every subgame. Every finite game has a subgame perfect equilibrium.

A.3 Self-Enforcing Behavior in Repeated Games: Subgame Perfect Equilibria, the Folk Theorem, and Imperfect Monitoring

So far we have examined games in which players interact only once. Institutional analysis, however, is concerned with recurrent situations, in which individuals interact over time. One way to examine such situations is to use dynamic games with more complicated game trees. A subset of such games—repeated games—has been found to be particularly amenable to formal analysis and useful for institutional analysis (Chapter 6).

Repeated-game theory examines situations in which the same (dynamic or static) stage game (such as a prisoners' dilemma or one-sided prisoner's dilemma game) is repeated every period. At the end of each period, payoffs are allocated, information might be revealed, and the same stage game is repeated again. Future payoffs are discounted by a time discount factor (often denoted by δ). A history in a repeated game is the set of actions taken in the past; a strategy specifies action combination in every stage game after every possible history. A strategy combination specifies a strategy for each player.[8]

To examine self-enforcing behavior in such games, suppose that the stage game is the prisoners' dilemma game presented in Figure A.1. If this stage game is repeated only once, the

[8] For ease of presentation, I often refer to an action combination as a strategy.

only subgame perfect equilibrium is (defect, defect); (cooperate, cooperate) is not an equilibrium. A comparable subgame perfect equilibrium in the repeated game is that after every history both players always defect. This equilibrium is also the unique equilibrium if the game is repeated a finite number of times. The reasons for this and the implied important implications for institutional analysis are discussed in Appendix C, section C.2.1. The discussion here focuses on situations in which the stage game is repeated for an infinite number of periods.

When the stage game is infinitely repeated, the preceding strategy is still a subgame perfect equilibrium. Each player's best response to this strategy is always to defect. But other equilibria are also possible.[9] Consider, for example, the following strategy to each player: In the first period, cooperate. Thereafter cooperate if all moves in all previous periods have been (cooperate, cooperate); otherwise defect. Each player's strategy thus calls for initiating exchange in the first period and cooperating as long as the other also cooperates. It calls for no cooperation if either player ever defects. This threat of ceasing cooperation forever is credible because (defect, defect) is an equilibrium.

A credible threat of such a trigger strategy can motivate the players to cooperate if they are sufficiently patient. The strategy implies that a player has to choose between present and future gains. Defection implies a relatively large immediate gain (5 in the game presented in Figure A.1), because the other player cooperates. But doing so implies losing future gains from cooperation because, following defection, both players will defect forever (and hence each will receive -8). The net present value of following the trigger strategy is $1/(1 - \delta)$. Deviating from it implies receiving a one-time payoff of 5, followed by -8 each period thereafter. This yields the net present value of $5 - 8/(1 - \delta)$, which declines as the players' time discount factor increases: if the players are sufficiently patient—if they value future gains enough—the preceding strategy is an equilibrium.

One of the most useful features of repeated-game theory is that verifying that a particular strategy combination is a subgame perfect equilibrium is often easier than verifying that a strategy is a Nash equilibrium. Roughly speaking, in any repeated game a strategy combination is a subgame perfect equilibrium if no player can gain from a one-period deviation after any history. In other words, to check if a particular strategy combination is a subgame perfect

[9] Experimental evidence indicates that people do indeed understand the strategic difference between one-shot and repeated games. See Appendix C.

equilibrium, it is sufficient to substantiate that after any history—any sequence of actions that can transpire given the strategy—no player can gain from a one-period deviation after which he will return to follow the strategy.[10]

In strategic dynamic situations, multiple equilibria often exist. The folk theorem of repeated games established that in infinitely repeated games there is usually an infinite number of subgame perfect equilibria.[11] Given the rules of the game, more than one pattern of behavior can prevail as an equilibrium outcome, and this is more likely to be the case in dynamic games with large actions set.

By revealing the general existence of multiple equilibria, game theory raises the problem of equilibrium selection. The "refinement" literature in game theory has attempted to refine the concept of the Nash equilibrium to restrict the set of admissible outcomes deductively. Subgame perfect equilibrium is one such restriction. But so far it has not offered a suitable deductive refinement for infinite repeated games (Van Damme 1983, 1987; Fudenberg and Tirole 1991).

[10] The formal analysis is due to Abreu (1988). Definition: Consider a strategy combination s, and denote the set of players by N and a player by i. The strategy is made up of s_i, the strategy for player i, and s_{-i}, the strategy for the other players. The strategy s_i is unimprovable against s_{-i} if there is no t - 1 period history (for any t) after which i could profit by deviating from s_i in period t only (and conforming to s_i from t + 1 and on). Proposition: Let the payoffs of a stage game G be bounded. In each finitely or infinitely repeated version of the game with time discount factor $\beta \in (0, 1)$, a strategy σ is a subgame perfect equilibrium if and only if for $\forall$ i (i.e., every player), σ_i is unimprovable against σ.

[11] The original folk theorem of repeated games (Friedman 1971) established that any average payoff vector that is better for all players than the (static, one-period game) Nash equilibrium payoff vector can be sustained as the outcome of a subgame perfect equilibrium of the infinitely repeated game if the players are sufficiently patient. Later analyses established that the equilibrium outcome set is even larger (see, e.g., Fudenberg and Maskin 1986).