

Chapter 2

GAME THEORY IN SUPPLY CHAIN ANALYSIS

G erard P. Cachon

The Wharton School

University of Pennsylvania

Philadelphia, PA 19104

cachon@wharton.upenn.edu

Serguei Netessine

The Wharton School

University of Pennsylvania

Philadelphia, PA 19104

netessine@wharton.upenn.edu

Abstract Game theory has become an essential tool in the analysis of supply chains with multiple agents, often with conflicting objectives. This chapter surveys the applications of game theory to supply chain analysis and outlines game-theoretic concepts that have potential for future application. We discuss both non-cooperative and cooperative game theory in static and dynamic settings. Careful attention is given to techniques for demonstrating the existence and uniqueness of equilibrium in non-cooperative games. A newsvendor game is employed throughout to demonstrate the application of various tools.¹

Keywords: Game theory, non-cooperative, cooperative, equilibrium concepts

1. Introduction

Game theory (hereafter GT) is a powerful tool for analyzing situations in which the decisions of multiple agents affect each agent's payoff. As such, GT deals with interactive optimization problems. While many economists in the past few centuries have worked on what can be considered game-theoretic models, John von Neumann and Oskar Morgenstern are formally credited as the fathers of modern game theory. Their classic book "Theory of Games and Economic Behavior", von Neumann and Morgenstern (1944), summarizes the basic concepts existing at that time. GT has since enjoyed an explosion of developments, including the concept of equilibrium by Nash (1950), games with imperfect information by Kuhn (1953), cooperative games by Aumann (1959) and Shubik (1962) and auctions by Vickrey (1961), to name just a few. Citing Shubik (2002), "In the 50s ... game theory was looked upon as a curiosity not to be taken seriously by any behavioral scientist. By the late 1980s, game theory in the new industrial organization has taken over ... game theory has proved its success in many disciplines."

This chapter has two goals. In our experience with GT problems we have found that many of the useful theoretical tools are spread over dozens of papers and books, buried among other tools that are not as useful in supply chain management (hereafter SCM). Hence, our first goal is to construct a brief tutorial through which SCM researchers can quickly locate GT tools and apply GT concepts. Due to the need for short explanations, we omit all proofs, choosing to focus only on the intuition behind the results we discuss. Our second goal is to provide ample but by no means exhaustive references on the specific applications of various GT techniques. These references offer an in-depth understanding of an application where necessary. Finally, we intentionally do not explore the implications of GT analysis on supply chain management, but rather, we emphasize the means of conducting the analysis to keep the exposition short.

1.1 Scope and relation to the literature

There are many GT concepts, but this chapter focuses on concepts that are particularly relevant to SCM and, perhaps, already found their applications in the literature. We dedicate a considerable amount of space to the discussion of static non-cooperative, non-zero sum games, the type of game which has received the most attention in the recent SCM literature. We also discuss cooperative games, dynamic/differential games and games with asymmetric/incomplete information. We omit

discussion of important GT concepts covered in other chapters in this book: auctions in Chapters 4 and 10; principal-agent models in Chapter 3; and bargaining in Chapter 11.

The material in this chapter was collected predominantly from Friedman (1986), Fudenberg and Tirole (1991), Moulin (1986), Myerson (1997), Topkis (1998) and Vives (1999). Some previous surveys of GT models in management science include Lucas's (1971) survey of mathematical theory of games, Feichtinger and Jorgensen's (1983) survey of differential games and Wang and Parlar's (1989) survey of static models. A recent survey by Li and Whang (2001) focuses on application of GT tools in five specific OR/MS models.

2. Non-cooperative static games

In non-cooperative static games the players choose strategies simultaneously and are thereafter committed to their chosen strategies, i.e., these are simultaneous move, one-shot games. Non-cooperative GT seeks a rational prediction of how the game will be played in practice.² The solution concept for these games was formally introduced by John Nash (1950) although some instances of using similar concepts date back a couple of centuries.

2.1 Game setup

To break the ground for the section, we introduce basic GT notation. A warning to the reader: to achieve brevity, we intentionally sacrifice some precision in our presentation. See texts like Friedman (1986) and Fudenberg and Tirole (1991) if more precision is required.

Throughout this chapter we represent games in the normal form. A game in the normal form consists of (1) *players* indexed by $i = 1, \dots, n$, (2) *strategies* or more generally a set of strategies denoted by x_i , $i = 1, \dots, n$ available to each player and (3) *payoffs* $\pi_i(x_1, x_2, \dots, x_n)$, $i = 1, \dots, n$ received by each player. Each strategy is defined on a set X_i , $x_i \in X_i$, so we call the Cartesian product $X_1 \times X_2 \times \dots \times X_n$ the *strategy space*. Each player may have a unidimensional strategy or a multi-dimensional strategy. In most SCM applications players have unidimensional strategies, so we shall either explicitly or implicitly assume unidimensional strategies throughout this chapter. Furthermore, with the exception of one example, we will work with continuous strategies, so the strategy space is R^n .

A player's strategy can be thought of as the complete instruction for which actions to take in a game. For example, a player can give his or her strategy to a person that has absolutely no knowledge of the player's payoff or preferences and that person should be able to use the instructions contained in the strategy to choose the actions the player desires. As a result, each player's set of feasible strategies must be independent of the strategies chosen by the other players, i.e., the strategy choice by one player is not allowed to limit the feasible strategies of another player. (Otherwise the game is ill defined and any analytical results obtained from the game are questionable.)

In the normal form players choose strategies simultaneously. Actions are adopted after strategies are chosen and those actions correspond to the chosen strategies.

As an alternative to the one-shot selection of strategies in the normal form, a game can also be designed in the extensive form. With the extensive form actions are chosen only as needed, so sequential choices are possible. As a result, players may learn information between the selection of actions, in particular, a player may learn which actions were previously chosen or the outcome of a random event. Figure 2.1 provides an example of a simple extensive form game and its equivalent normal form representation: there are two players, player I chooses from {Left, Right} and player II chooses from {Up, Down}. In the extensive form player I chooses first, then player II chooses after learning player I's choice. In the normal form they choose simultaneously. The key distinction between normal and extensive form games is that in the normal form a player is able to commit to all future decisions. We later show that this additional commitment power may influence the set of plausible equilibria.

A player can choose a particular strategy or a player can choose to randomly select from among a set of strategies. In the former case the player is said to choose a *pure strategy* whereas in the latter case the player chooses a *mixed strategy*. There are situations in economics and marketing that have used mixed strategies: see, e.g., Varian (1980) for search models and Lal (1990) for promotion models. However, mixed strategies have not been applied in SCM, in part because it is not clear how a manager would actually implement a mixed strategy. For example, it seems unreasonable to suggest that a manager should "flip a coin" among various capacity levels. Fortunately, mixed strategy equilibria do not exist in games with a unique pure strategy equilibrium. Hence,

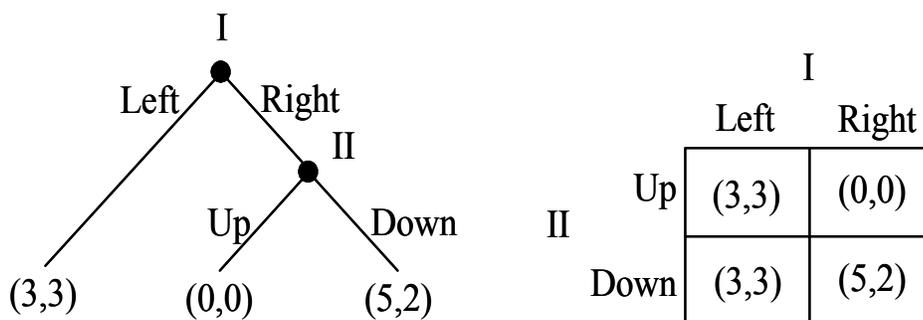


Figure 2.1. Extensive vs normal form game representation.

in those games attention can be restricted to pure strategies without loss of generality. Therefore, in the remainder of this chapter we consider only pure strategies.

In a *non-cooperative* game the players are unable to make binding commitments before choosing their strategies. In a *cooperative* game players are able to make binding commitments. Hence, in a cooperative game players can make side-payments and form coalitions. We begin our analysis with non-cooperative static games.

In all sections, except the last one, we work with the games of *complete information*, i.e., the players' strategies and payoffs are common knowledge to all players.

As a practical example throughout this chapter, we utilize the classic newsvendor problem transformed into a game. In the absence of competition each newsvendor buys Q units of a single product at the beginning of a single selling season. Demand during the season is a random variable D with distribution function F_D and density function f_D . Each unit is purchased for c and sold on the market for $r > c$. The newsvendor solves the following optimization problem

$$\max_Q \pi = \max_Q E_D [r \min(D, Q) - cQ],$$

with the unique solution

$$Q^* = F_D^{-1} \left(\frac{r - c}{r} \right).$$

Goodwill penalty costs and salvage revenues can easily be incorporated into the analysis, but for our needs we normalized them out.

Now consider the GT version of the newsvendor problem with two retailers competing on product availability. Parlar (1988) was the first to analyze this problem, which is also one of the first articles modeling inventory management in a GT framework. It is useful to consider only the two-player version of this game because then graphical analysis and interpretations are feasible. Denote the two players by subscripts i and j , their strategies (in this case stocking quantities) by Q_i, Q_j and their payoffs by π_i, π_j .

We introduce interdependence of the players' payoffs by assuming the two newsvendors sell the same product. As a result, if retailer i is out of stock, all unsatisfied customers try to buy the product at retailer j instead. Hence, retailer i 's total demand is $D_i + (D_j - Q_j)^+$: the sum of his own demand and the demand from customers not satisfied by retailer j . Payoffs to the two players are then

$$\pi_i(Q_i, Q_j) = E_D \left[r_i \min \left(D_i + (D_j - Q_j)^+, Q_i \right) - c_i Q_i \right], \quad i, j = 1, 2.$$

2.2 Best response functions and the equilibrium of the game

We are ready for the first important GT concept: *best response functions*.

DEFINITION 1. *Given an n -player game, player i 's best response (function) to the strategies x_{-i} of the other players is the strategy x_i^* that maximizes player i 's payoff $\pi_i(x_i, x_{-i})$:*

$$x_i^*(x_{-i}) = \arg \max_{x_i} \pi_i(x_i, x_{-i}).$$

($x_i^*(x_{-i})$ is probably better described as a correspondence rather than a function, but we shall nevertheless call it a function with an understanding that we are interpreting the term "function" liberally.) If π_i is quasi-concave in x_i the best response is uniquely defined by the first-order conditions of the payoff functions. In the context of our competing newsvendors example, the best response functions can be found by optimizing each player's payoff functions w.r.t. the player's own decision variable Q_i while taking the competitor's strategy Q_j as given. The

resulting best response functions are

$$Q_i^*(Q_j) = F_{D_i+(D_j-Q_j)^+}^{-1} \left(\frac{r_i - c_i}{r_i} \right), \quad i, j = 1, 2.$$

Taken together, the two best response functions form a *best response mapping* $R^2 \rightarrow R^2$ or in the more general case $R^n \rightarrow R^n$. Clearly, the best response is the best player i can hope for given the decisions of other players. Naturally, an outcome in which all players choose their best responses is a candidate for the non-cooperative solution. Such an outcome is called a Nash equilibrium (hereafter NE) of the game.

DEFINITION 2. *An outcome $(x_1^*, x_2^*, \dots, x_n^*)$ is a Nash equilibrium of the game if x_i^* is a best response to x_{-i}^* for all $i = 1, 2, \dots, n$.*

Going back to competing newsvendors, NE is characterized by solving a *system* of best responses that translates into the system of first-order conditions:

$$\begin{aligned} Q_1^*(Q_2^*) &= F_{D_1+(D_2-Q_2^*)^+}^{-1} \left(\frac{r_1 - c_1}{r_1} \right), \\ Q_2^*(Q_1^*) &= F_{D_2+(D_1-Q_1^*)^+}^{-1} \left(\frac{r_2 - c_2}{r_2} \right). \end{aligned}$$

When analyzing games with two players it is often helpful to graph the best response functions to gain intuition. Best responses are typically defined implicitly through the first-order conditions, which makes analysis difficult. Nevertheless, we can gain intuition by finding out how each player reacts to an increase in the stocking quantity by the other player (i.e., $\partial Q_i^*(Q_j) / \partial Q_j$) through employing implicit differentiation as follows:

$$\frac{\partial Q_i^*(Q_j)}{\partial Q_j} = - \frac{\frac{\partial^2 \pi_i}{\partial Q_i \partial Q_j}}{\frac{\partial^2 \pi_i}{\partial Q_i^2}} = - \frac{r_i f_{D_i+(D_j-Q_j)^+ | D_j > Q_j} (Q_i) \Pr (D_j > Q_j)}{r_i f_{D_i+(D_j-Q_j)^+} (Q_i)} < 0. \quad (2.1)$$

The expression says that the *slopes* of the best response functions are negative, which implies an intuitive result that each player's best response is monotonically decreasing in the other player's strategy. Figure 2.2 presents this result for the symmetric newsvendor game. The equilibrium is located on the intersection of the best responses and we also see that the best responses are, indeed, decreasing.

One way to think about a NE is as a *fixed point* of the best response mapping $R^n \rightarrow R^n$. Indeed, according to the definition, NE must

satisfy the system of equations $\partial\pi_i/\partial x_i = 0$, all i . Recall that a fixed point x of mapping $f(x)$, $R^n \rightarrow R^n$ is any x such that $f(x) = x$. Define $f_i(x_1, \dots, x_n) = \partial\pi_i/\partial x_i + x_i$. By the definition of a fixed point,

$$f_i(x_1^*, \dots, x_n^*) = x_i^* = \partial\pi_i(x_1^*, \dots, x_n^*)/\partial x_i + x_i^* \rightarrow \partial\pi_i(x_1^*, \dots, x_n^*)/\partial x_i = 0, \text{ all } i.$$

Hence, x^* solves the first-order conditions if and only if it is a fixed point of mapping $f(x)$ defined above.

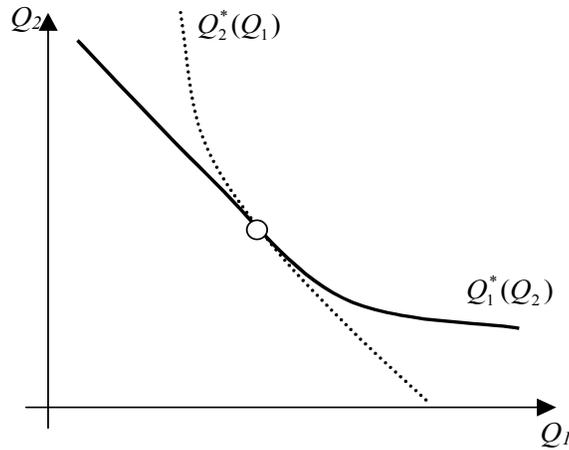


Figure 2.2. Best responses in the newsvendor game.

The concept of NE is intuitively appealing. Indeed, it is a self-fulfilling prophecy. To explain, suppose a player were to guess the strategies of the other players. A guess would be consistent with payoff maximization and therefore would be reasonable only if it presumes that strategies are chosen to maximize every player's payoff given the chosen strategies. In other words, with any set of strategies that is not a NE there exists at least one player that is choosing a non payoff maximizing strategy. Moreover, the NE has a self-enforcing property: no player wants to unilaterally deviate from it since such behavior would lead to lower payoffs. Hence NE seems to be the necessary condition for the prediction of any rational behavior by players³.

While attractive, numerous criticisms of the NE concept exist. Two particularly vexing problems are the non-existence of equilibrium and the multiplicity of equilibria. Without the existence of an equilibrium, little can be said regarding the likely outcome of the game. If there are multiple equilibria, then it is not clear which one will be the outcome.

Indeed, it is possible the outcome is not even an equilibrium because the players may choose strategies from different equilibria. For example, consider the normal form game in Figure 2.1. There are two Nash equilibria in that game $\{\text{Left,Up}\}$ and $\{\text{Right,Down}\}$: each is a best response to the other player's strategy. However, because the players choose their strategies simultaneously it is possible that player I chooses Right (the 2nd equilibrium) while player II choose Up (the 1st equilibrium), which results in $\{\text{Right,Up}\}$, the worst outcome for both players.

In some situations it is possible to rationalize away some equilibria via a refinement of the NE concept: e.g., trembling hand perfect equilibrium by Selten (1975), sequential equilibrium by Kreps and Wilson (1982) and proper equilibria by Myerson (1997). These refinements eliminate equilibria that are based on non-credible threats, i.e., threats of future actions that would not actually be adopted if the sequence of events in the game led to a point in the game in which those actions could be taken. The extensive form game in Figure 2.1 illustrates this point. $\{\text{Left, Up}\}$ is a Nash equilibrium (just as it is in the comparable normal form game) because each player is choosing a best response to the other player's strategy: Left is optimal for player I given player II plans to play Up and player II is indifferent between Up or Down given player I chooses Left. But if player I were to choose Right, then it is unreasonable to assume player II would actually follow through with UP: UP yields a payoff of 0 while Down yields a payoff of 2. Hence, the $\{\text{Left, Up}\}$ equilibrium is supported by a non-credible threat by player II to play Up. Although these refinements are viewed as extremely important in economics (Selten was awarded the Nobel prize for his work), the need for these refinements has not yet materialized in the SCM literature. But that may change as more work is done on sequential/dynamic games.

An interesting feature of the NE concept is that the system optimal solution (i.e., a solution that maximizes the sum of players' payoffs) need not be a NE. Hence, decentralized decision making generally introduces inefficiency in the supply chain. There are, however, some exceptions: see Mahajan and van Ryzin (1999b) and Netessine and Zhang (2003) for situations in which competition may result in the system-optimal performance. In fact, a NE may not even be on the *Pareto frontier*: the set of strategies such that each player can be made better off only if some other player is made worse off. A set of strategies is *Pareto optimal* if they are on the Pareto frontier; otherwise a set of strategies is *Pareto inferior*. Hence, a NE can be Pareto inferior. The Prisoner's Dilemma

game (see Fudenberg and Tirole 1991) is the classic example of this: only one pair of strategies when both players “cooperate” is Pareto optimal, and the unique Nash equilibrium is when both players “defect” happens to be Pareto inferior. A large body of the SCM literature deals with ways to align the incentives of competitors to achieve optimality. See Cachon (2002) for a comprehensive survey and taxonomy. See Cachon (2003) for a supply chain analysis that makes extensive use of the Pareto optimal concept.

2.3 Existence of equilibrium

A NE is a solution to a system of n first-order conditions, so an equilibrium may not exist. Non-existence of an equilibrium is potentially a conceptual problem since in this case it is not clear what the outcome of the game will be. However, in many games a NE does exist and there are some reasonably simple ways to show that at least one NE exists. As already mentioned, a NE is a fixed point of the best response mapping. Hence fixed point theorems can be used to establish the existence of an equilibrium. There are three key fixed point theorems, named after their creators: Brouwer, Kakutani and Tarski, see Border (1999) for details and references. However, direct application of fixed point theorems is somewhat inconvenient and hence generally not done. For exceptions see Lederer and Li (1997) and Majumder and Groenevelt (2001a) for existence proofs that are based on Brouwer’s fixed point theorem. Alternative methods, derived from these fixed point theorems, have been developed. The simplest and the most widely used technique for demonstrating the existence of NE is through verifying concavity of the players’ payoffs.

THEOREM 1 (Debreu 1952). *Suppose that for each player the strategy space is compact⁴ and convex and the payoff function is continuous and quasi-concave with respect to each player’s own strategy. Then there exists at least one pure strategy NE in the game.*

If the game is symmetric in a sense that the players’ strategies and payoffs are identical, one would imagine that a symmetric solution should exist. This is indeed the case, as the next Theorem ascertains.

THEOREM 2. *Suppose that a game is symmetric and for each player the strategy space is compact and convex and the payoff function is continuous and quasi-concave with respect to each player’s own strategy.*

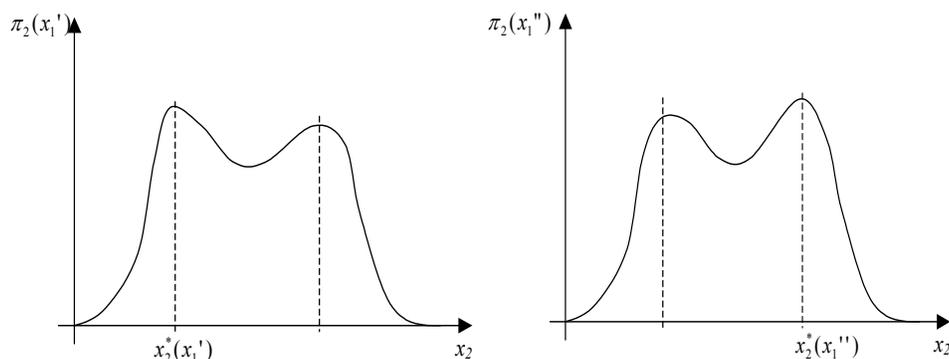


Figure 2.3. Example with a bi-modal objective function.

Then there exists at least one symmetric pure strategy NE in the game.

To gain some intuition about why non-quasi-concave payoffs may lead to non-existence of NE, suppose that in a two-player game, player 2 has a bi-modal objective function with two local maxima. Furthermore, suppose that a small change in the strategy of player 1 leads to a shift of the global maximum for player 2 from one local maximum to another. To be more specific, let us say that at x_1' the global maximum $x_2^*(x_1')$ is on the left (Figure 2.3 left) and at x_1'' the global maximum $x_2^*(x_1'')$ is on the right (Figure 2.3 right). Hence, a small change in x_1 from x_1' to x_1'' induces a jump in the best response of player 2, x_2^* . The resulting best response mapping is presented in Figure 2.4 and there is no NE in pure strategies in this game. In other words, best response functions do not intersect anywhere. As a more specific example, see Netessine and Shumsky (2001) for an extension of the newsvendor game to the situation in which product inventory is sold at two different prices; such a game may not have a NE since both players' objectives may be bimodal. Furthermore, Cachon and Harker (2002) demonstrate that pure strategy NE may not exist in two other important settings: two retailers competing with cost functions described by the Economic Order Quantity (EOQ) or two service providers competing with service times described by the $M/M/1$ queuing model.

The assumption of a compact strategy space may seem restrictive. For example, in the newsvendor game the strategy space R_+^2 is not bounded from above. However, we could easily bound it with some large enough finite number to represent the upper bound on the demand distribution. That bound would not impact any of the choices,

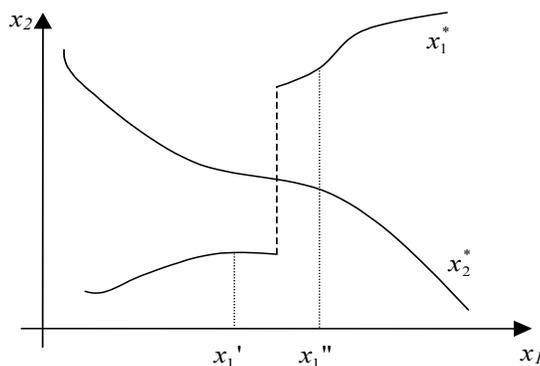


Figure 2.4. Non-existence of NE.

and therefore the transformed game behaves just as the original game with an unbounded strategy space. (However, that bound cannot depend on any player's strategy choice.)

To continue with the newsvendor game analysis, it is easy to verify that the newsvendor's objective function is concave and hence quasi-concave w.r.t. the stocking quantity by taking the second derivative. Hence the conditions of Theorem 1 are satisfied and a NE exists. There are virtually dozens of papers employing Theorem 1. See, for example, Lippman and McCardle (1997) for the proof involving quasi-concavity, Mahajan and van Ryzin (1999a) and Netessine et al. (2002) for the proofs involving concavity. Clearly, quasi-concavity of each player's objective function only implies uniqueness of the best response but does not imply a unique NE. One can easily envision a situation where unique best response functions cross more than once so that there are multiple equilibria (see Figure 2.5).

If quasi-concavity of the players' payoffs cannot be verified, there is an alternative existence proof that relies on Tarski's (1955) fixed point theorem and involves the notion of supermodular games. The theory of supermodular games is a relatively recent development introduced and advanced by Topkis (1998).

DEFINITION 3. *A twice continuously differentiable payoff function $\pi_i(x_1, \dots, x_n)$ is supermodular (submodular) iff $\partial^2 \pi_i / \partial x_i \partial x_j \geq 0$ (≤ 0) for all x and all $j \neq i$. The game is called supermodular if the players' payoffs are supermodular.*

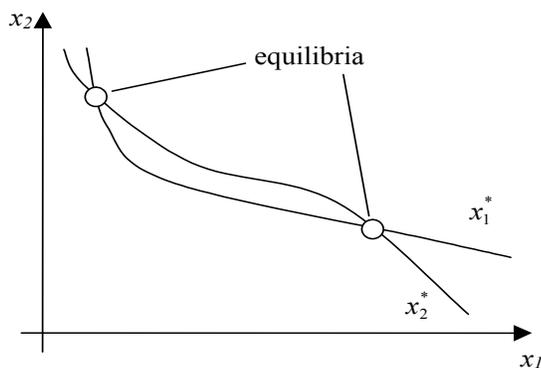


Figure 2.5. Non-uniqueness of the equilibrium.

Supermodularity essentially means complementarity between any two strategies and is not linked directly to either convexity, concavity or even continuity. (This is a significant advantage when forced to work with discrete strategies, e.g., Cachon 2001.) However, similar to concavity/convexity, supermodularity/submodularity is preserved under maximization, limits and addition and hence under expectation/integration signs, an important feature in stochastic SCM models. While in most situations the positive sign of the second derivative can be used to verify supermodularity (using Definition 3), sometimes it is necessary to utilize supermodularity-preserving transformations to show that payoffs are supermodular. Topkis (1998) provides a variety of ways to verify that the function is supermodular and some of his results are used in Cachon (2001), Cachon and Lariviere (1999), Netessine and Shumsky (2001) and Netessine and Rudi (2003). The following theorem follows directly from Tarski’s fixed point result and provides another tool to show existence of NE in non-cooperative games:

THEOREM 3. *In a supermodular game there exists at least one NE.*

Coming back to the competitive newsvendors example, recall that the second-order cross-partial derivative was found to be

$$\frac{\partial^2 \pi_i}{\partial Q_i \partial Q_j} = -r_i f_{D_i+(D_j-Q_j)+|D_j>Q_j}(Q_i) \Pr(D_j > Q_j) < 0,$$

so that the newsvendor game is submodular and hence existence of equilibrium cannot be assured. However, a standard trick is to re-define the

ordering of the players' strategies. Let $y = -Q_j$ so that

$$\frac{\partial^2 \pi_i}{\partial Q_i \partial y} = r_i f_{D_i + (D_j + y) + |D_j > Q_j} (Q_i) \Pr (D_j > -y) > 0,$$

and the game becomes supermodular in (x_i, y) so existence of NE is assured. Notice that we do not change either payoffs or the structure of the game, we only alter the ordering of one player's strategy space. Obviously, this trick only works in two-player games, see also Lippman and McCardle (1997) for the analysis of the more general version of the newsvendor game using a similar transformation. Hence, we can state that in general NE exists in games with decreasing best responses (submodular games) with two players. This argument can be generalized slightly in two ways that we mention briefly, see Vives (1999) for details. One way is to consider an n -player game where best responses are functions of aggregate actions of all other players, that is, $x_i^* = x_i^* \left(\sum_{j \neq i} x_j \right)$. If best responses in such a game are decreasing, then NE exists. Another generalization is to consider the same game with $x_i^* = x_i^* \left(\sum_{j \neq i} x_j \right)$ but require symmetry. In such a game, existence can be shown even with non-monotone best responses provided that there are only jumps up but on intervals between jumps best responses can be increasing or decreasing.

We now step back to discuss the intuition behind the supermodularity results. Roughly speaking, Tarski's fixed point theorem only requires best response mappings to be non-decreasing for the existence of equilibrium and does not require quasi-concavity of the players' payoffs and allows for jumps in best responses. While it may be hard to believe that non-decreasing best responses is the only requirement for the existence of a NE, consider once again the simplest form of a single-dimensional equilibrium as a solution to the fixed point mapping $x = f(x)$ on the compact set. It is easy to verify after a few attempts that if $f(x)$ is non-decreasing but possibly with jumps up then it is not possible to derive a situation without an equilibrium. However, when $f(x)$ jumps down, non-existence is possible (see Figure 2.6).

Hence, increasing best response functions is the only major requirement for an equilibrium to exist; players' objectives do not have to be quasi-concave or even continuous. However, to describe an existence theorem with non-continuous payoffs requires the introduction of terms and definitions from lattice theory. As a result, we restricted ourselves to the assumption of continuous payoff functions, and in particular, to twice-differentiable payoff functions.

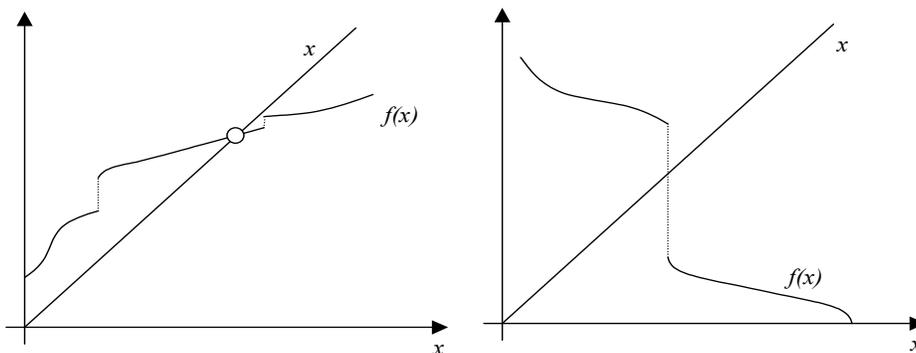


Figure 2.6. Increasing (left) and decreasing (right) mappings.

Although it is now clear why increasing best responses ensure existence of an equilibrium, it is not immediately obvious why Definition 3 provides a sufficient condition, given that it only concerns the sign of the second-order cross-partial derivative. To see this connection, consider separately the continuous and the dis-continuous parts of the best response $x_i^*(x_j)$. When the best response is continuous, we can apply the Implicit Function Theorem to find its slope as follows

$$\frac{\partial x_i^*}{\partial x_j} = -\frac{\frac{\partial^2 \pi_i}{\partial x_i \partial x_j}}{\frac{\partial^2 \pi_i}{\partial x_i^2}}.$$

Clearly, if x_i^* is the best response, it must be the case that $\partial^2 \pi_i / \partial x_i^2 < 0$ or else it would not be the best response. Hence, for the slope to be positive it is sufficient to have $\partial^2 \pi_i / \partial x_i \partial x_j > 0$ which is what Definition 3 provides. This reasoning does not, however, work at discontinuities in best responses since the Implicit Function Theorem cannot be applied. To show that only jumps up are possible if $\partial^2 \pi_i / \partial x_i \partial x_j > 0$ holds, consider a situation in which there is a jump down in the best response. As one can recall, jumps in best responses happen when the objective function is bi-modal (or more generally multi-modal). For example, consider a specific point $x_j^\#$ and let $x_i^1(x_j^\#) < x_i^2(x_j^\#)$ be two distinct points at which first-order conditions hold (i.e., the objective function π_i is bi-modal). Further, suppose $\pi_i(x_i^1(x_j^\#), x_j^\#) < \pi_i(x_i^2(x_j^\#), x_j^\#)$ but $\pi_i(x_i^1(x_j^\# + \varepsilon), x_j^\# + \varepsilon) > \pi_i(x_i^2(x_j^\# + \varepsilon), x_j^\# + \varepsilon)$. That is, initially $x_i^2(x_j^\#)$ is a global maximum but as we increase $x_j^\#$ infinitesimally, there

is a *jump down* and a smaller $x_i^1(x_j^\# + \varepsilon)$ becomes the global maximum. For this to be the case, it must be that

$$\frac{\partial \pi_i(x_i^1(x_j^\#), x_j^\#)}{\partial x_j} > \frac{\partial \pi_i(x_i^2(x_j^\#), x_j^\#)}{\partial x_j},$$

or, in words, the objective function rises faster at $(x_i^1(x_j^\#), x_j^\#)$ than at $(x_i^2(x_j^\#), x_j^\#)$. This, however, can only happen if $\partial^2 \pi_i / \partial x_i \partial x_j < 0$ at least somewhere on the interval $[x_i^1(x_j^\#), x_i^2(x_j^\#)]$ which is a contradiction. Hence, if $\partial^2 \pi_i / \partial x_i \partial x_j > 0$ holds then only jumps up in the best response are possible.

2.4 Uniqueness of equilibrium

From the perspective of generating qualitative insights, it is quite useful to have a game with a unique NE. If there is only one equilibrium, then one can characterize equilibrium actions without much ambiguity. Unfortunately, demonstrating uniqueness is generally much harder than demonstrating existence of equilibrium. This section provides several methods for proving uniqueness. No single method dominates; all may have to be tried to find the one that works. Furthermore, one should be careful to recognize that these methods assume existence, i.e., existence of NE must be shown separately. Finally, it is worth pointing out that uniqueness results are only available for games with continuous best response functions and hence there are no general methods to prove uniqueness of NE in supermodular games.

2.4.1 Method 1. Algebraic argument. In some rather fortunate situations one can ascertain that the solution is unique by simply looking at the optimality conditions. For example, in a two-player game the optimality condition of one of the players may have a unique closed-form solution that does not depend on the other player's strategy and, given the solution for one player, the optimality condition for the second player can be solved uniquely. See Hall and Porteus (2000) and Netessine and Rudi (2001) for examples. In other cases one can assure uniqueness by analyzing geometrical properties of the best response functions and arguing that they intersect only once. Of course, this is only feasible in two-player games. See Parlar (1988) for a proof of uniqueness in the two-player newsvendor game and Majumder and Groenevelt (2001b) for a supply chain game with competition in reverse logistics. However, in most situations these geometrical properties are

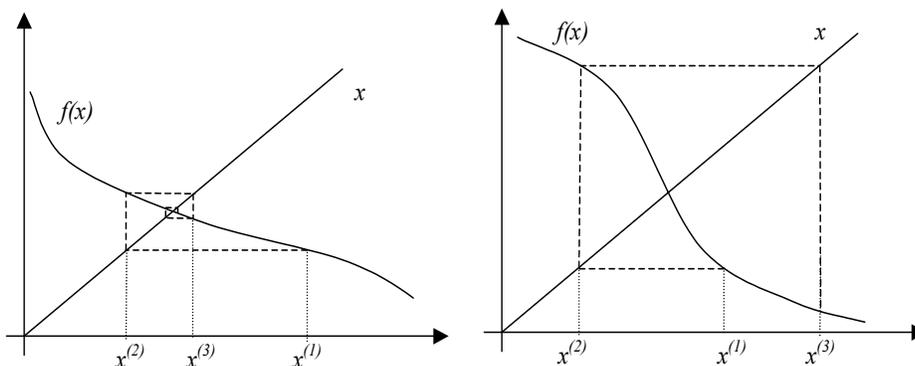


Figure 2.7. Converging (left) and diverging (right) iterations.

also implied by the more formal arguments stated below. Finally, it may be possible to use a contradiction argument: assume that there is more than one equilibrium and prove that such an assumption leads to a contradiction, as in Lederer and Li (1997).

2.4.2 Method 2. Contraction mapping argument.

Although the most restrictive among all methods, the contraction mapping argument is the most widely known and is the most frequently used in the literature because it is the easiest to verify. The argument is based on showing that the best response mapping is a contraction, which then implies the mapping has a unique fixed point. To illustrate the concept of a contraction mapping, suppose we would like to find a solution to the following fixed point equation:

$$x = f(x), x \in R^1.$$

To do so, a sequence of values is generated by an iterative algorithm, $\{x^{(1)}, x^{(2)}, x^{(3)}, \dots\}$ where $x^{(1)}$ is arbitrarily picked and $x^{(t)} = f(x^{(t-1)})$. The hope is that this sequence converges to a unique fixed point. It does so if, roughly speaking, each step in the sequence moves closer to the fixed point. One could verify that if $|f'(x)| < 1$ in some vicinity of x^* then such an iterative algorithm converges to a unique $x^* = f(x^*)$. Otherwise, the algorithm diverges. Graphically, the equilibrium point is located on the intersection of two functions: x and $f(x)$. The iterative algorithm is presented in Figure 2.7.

The iterative scheme in Figure 2.7 left is a contraction mapping: it approaches the equilibrium after every iteration.

DEFINITION 4. Mapping $f(x)$, $R^n \rightarrow R^n$ is a contraction iff $\|f(x_1) - f(x_2)\| \leq \alpha \|x_1 - x_2\|$, $\forall x_1, x_2, \alpha < 1$.

In words, the application of a contraction mapping to any two points strictly reduces (i.e., $\alpha = 1$ does not work) the distance between these points. The norm in the definition can be any norm, i.e., the mapping can be a contraction in one norm and not a contraction in another norm.

THEOREM 4. *If the best response mapping is a contraction on the entire strategy space, there is a unique NE in the game.*

One can think of a contraction mapping in terms of iterative play: player 1 selects some strategy, then player 2 selects a strategy based on the decision by player 1, etc. If the best response mapping is a contraction, the NE obtained as a result of such iterative play is *stable* but the opposite is not necessarily true, i.e., no matter where the game starts, the final outcome is the same. See also Moulin (1986) for an extensive treatment of stable equilibria.

A major restriction in Theorem 4 is that the contraction mapping condition must be satisfied everywhere. This assumption is quite restrictive because the best response mapping may be a contraction locally, say in some not necessarily small ε -neighborhood of the equilibrium but not outside of it. Hence, if iterative play starts in this ε -neighborhood, then it converges to the equilibrium, but starting outside that neighborhood may not lead to the equilibrium (even if the equilibrium is unique). Even though one may wish to argue that it is reasonable for the players to start iterative play some place close to the equilibrium, formalization of such an argument is rather difficult. Hence, we must impose the condition that the entire strategy space be considered. See Stidham (1992) for an interesting discussion of stability issues in a queuing system.

While Theorem 4 is a starting point towards a method for demonstrating uniqueness, it does not actually explain how to validate that a best reply mapping is a contraction. Suppose we have a game with n players each endowed with the strategy x_i and we have obtained the best response functions for all players, $x_i = f_i(x_{-i})$. We can then define the following matrix of derivatives of the best response functions:

$$A = \begin{vmatrix} 0 & \frac{\partial f_1}{\partial x_2} & \dots & \frac{\partial f_1}{\partial x_n} \\ \frac{\partial f_2}{\partial x_1} & 0 & \dots & \frac{\partial f_2}{\partial x_n} \\ \dots & \dots & \dots & \dots \\ \frac{\partial f_n}{\partial x_1} & \frac{\partial f_n}{\partial x_2} & \dots & 0 \end{vmatrix}.$$

Further, denote by $\rho(A)$ the spectral radius of matrix A and recall that the spectral radius of a matrix is equal to the largest absolute eigenvalue $\rho(A) = \{\max |\lambda| : Ax = \lambda x, x \neq 0\}$, see Horn and Johnson (1996).

THEOREM 5. *The mapping $f(x), R^n \rightarrow R^n$ is a contraction if and only if $\rho(A) < 1$ everywhere.*

Theorem 5 is simply an extension of the iterative convergence argument we used above into multiple dimensions, and the spectral radius rule is an extension of the requirement $|f'(x)| < 1$. Still, Theorem 5 is not as useful as we would like it to be: calculating eigenvalues of a matrix is not trivial. Instead, it is often helpful to use the fact that the largest eigenvalue and hence the spectral radius is bounded above by any of the matrix norms, see Horn and Johnson (1996). So instead of working with the spectral radius itself, it is sufficient to show $\|A\| < 1$ for any one matrix norm. The most convenient matrix norms are the maximum column-sum and the maximum row-sum norms (see Horn and Johnson 1996 for other matrix norms). To use either of these norms to verify the contraction mapping, it is sufficient to verify that no column sum or no row sum of matrix A exceeds one,

$$\sum_{i=1}^n \left| \frac{\partial f_k}{\partial x_i} \right| < 1, \text{ or } \sum_{i=1}^n \left| \frac{\partial f_i}{\partial x_k} \right| < 1, \forall k.$$

Netessine and Rudi (2003) used the contraction mapping argument in this most general form in the multiple-player variant of the newsvendor game described above.

A challenge associated with the contraction mapping argument is finding best response functions because in most SC models best responses cannot be found explicitly. Fortunately, Theorem 5 only requires the derivatives of the best response functions, which can be done using the Implicit Function Theorem (from now on, IFT, see Bertsekas 1999). Using the IFT, Theorem 5 can be re-stated as

$$\sum_{i=1, i \neq k}^n \left| \frac{\partial^2 \pi_k}{\partial x_k \partial x_i} \right| < \left| \frac{\partial^2 \pi_k}{\partial x_k^2} \right|, \forall k. \quad (2.2)$$

This condition is also known as “diagonal dominance” because the diagonal of the matrix of second derivatives, also called the Hessian, dominates

the off-diagonal entries:

$$H = \begin{vmatrix} \frac{\partial^2 \pi_1}{\partial x_1^2} & \frac{\partial^2 \pi_1}{\partial x_1 \partial x_2} & \cdots & \frac{\partial^2 \pi_1}{\partial x_1 \partial x_n} \\ \frac{\partial^2 \pi_2}{\partial x_2 \partial x_1} & \frac{\partial^2 \pi_2}{\partial x_2^2} & \cdots & \frac{\partial^2 \pi_2}{\partial x_2 \partial x_n} \\ \cdots & \cdots & \cdots & \cdots \\ \frac{\partial^2 \pi_n}{\partial x_n \partial x_1} & \frac{\partial^2 \pi_n}{\partial x_n \partial x_2} & \cdots & \frac{\partial^2 \pi_n}{\partial x_n^2} \end{vmatrix}. \quad (2.3)$$

Contraction mapping conditions in the diagonal dominance form have been used extensively by Bernstein and Federgruen (2000, 2001a, 2001c, 2003). As has been noted by Bernstein and Federgruen (2002), many standard economic demand models satisfy this condition.

In games with only two players the condition in Theorem 5 simplifies to

$$\left| \frac{\partial f_1}{\partial x_2} \right| < 1 \text{ and } \left| \frac{\partial f_2}{\partial x_1} \right| < 1, \quad (2.4)$$

i.e., the slopes of the best response functions are less than one. This condition is especially intuitive if we use the graphical illustration (Figure 2.2). Given that the slope of each best response function is less than one everywhere, if they cross at one point then they cannot cross at an additional point. A contraction mapping argument in this form was used by van Mieghem (1999) and by Rudi et al. (2001).

Returning to the newsvendor game example, we have found that the slopes of the best response functions are

$$\left| \frac{\partial Q_i^*(Q_j)}{\partial Q_j} \right| = \left| \frac{f_{D_i+(D_j-Q_j)^+|D_j>Q_j}(Q_i) \Pr(D_j > Q_j)}{f_{D_i+(D_j-Q_j)^+}(Q_i)} \right| < 1.$$

Hence, the best response mapping in the newsvendor game is a contraction and the game has a unique and stable NE.

2.4.3 Method 3. Univalent mapping argument. Another method for demonstrating uniqueness of equilibrium is based on verifying that the best response mapping is one-to-one: that is, if $f(x)$ is a $R^n \rightarrow R^n$ mapping, then $y = f(x)$ implies that for all $x' \neq x$, $y \neq f(x')$. Clearly, if the best response mapping is one-to-one then there can be at most one fixed point of such mapping. To make an analogy, recall that, if the equilibrium is interior⁵, the NE is a solution to the system of the first-order conditions: $\partial \pi_i / \partial x_i = 0$, $\forall i$, which defines the best response mapping. If this mapping is single-dimensional $R^1 \rightarrow R^1$ then it is quite clear that the condition sufficient for the mapping to be one-to-one is

quasi-concavity of π_i . Similarly, for the $R^n \rightarrow R^n$ mapping to be one-to-one we require quasi-concavity of the mapping which translates into quasi-definiteness of the Hessian:

THEOREM 6. *Suppose the strategy space of the game is convex and all equilibria are interior. Then if the determinant $|H|$ is negative quasi-definite (i.e., if the matrix $H + H^T$ is negative definite) on the players' strategy set, there is a unique NE.*

Proof of this result can be found in Gale and Nikaido (1965) and some further developments that deal with boundary equilibria are found in Rosen (1965). Notice that the univalent mapping argument is somewhat weaker than the contraction mapping argument. Indeed, the restatement (2.2) of the contraction mapping theorem directly implies univalence since the dominant diagonal assures us that H is negative definite. Hence, it is negative quasi-definite. It immediately follows that the newsvendor game satisfies the univalence theorem. However, if some other matrix norm is used, the relationship between the two theorems is not that specific. In the case of just two players the univalence theorem can be written as, according to Moulin (1986),

$$\left| \frac{\partial^2 \pi_2}{\partial x_2 \partial x_1} + \frac{\partial^2 \pi_1}{\partial x_1 \partial x_2} \right| \leq 2 \sqrt{\left| \frac{\partial^2 \pi_1}{\partial x_1^2} \cdot \frac{\partial^2 \pi_2}{\partial x_2^2} \right|}, \forall x_1, x_2.$$

2.4.4 Method 4. Index theory approach.. This method is based on the Poincare-Hopf index theorem found in differential topology, see, e.g., Gillemain and Pollack (1974). Similarly to the univalence mapping approach, it requires a certain sign from the Hessian, but this requirement need hold only at the equilibrium point.

THEOREM 7. *Suppose the strategy space of the game is convex and all payoff functions are quasi-concave. Then if $(-1)^n |H|$ is positive whenever $\partial \pi_i / \partial x_i = 0$, all i , there is a unique NE.*

Observe that the condition $(-1)^n |H|$ is trivially satisfied if $|H|$ is negative definite which is implied by the condition (2.2) of contraction mapping, i.e., this method is also somewhat weaker than the contraction mapping argument. Moreover, the index theory condition need only hold at the equilibrium. This makes it the most general, but also the hardest to apply. To gain some intuition about why the index theory method works, consider the two-player game. The condition of Theorem

7 simplifies to

$$\left| \begin{array}{cc} \frac{\partial^2 \pi_1}{\partial x_1^2} & \frac{\partial^2 \pi_1}{\partial x_1 \partial x_2} \\ \frac{\partial^2 \pi_2}{\partial x_2 \partial x_1} & \frac{\partial^2 \pi_2}{\partial x_2^2} \end{array} \right| > 0 \quad \forall x_1, x_2 : \frac{\partial \pi_1}{\partial x_1} = 0, \frac{\partial \pi_2}{\partial x_2} = 0,$$

which can be interpreted as meaning the multiplication of the slopes of best response functions should not exceed one at the equilibrium:

$$\frac{\partial f_1}{\partial x_2} \frac{\partial f_2}{\partial x_1} < 1 \text{ at } x_1^*, x_2^*. \quad (2.5)$$

As with the contraction mapping approach, with two players the Theorem becomes easy to visualize. Suppose we have found best response functions $x_1^* = f_1(x_2)$ and $x_2^* = f_2(x_1)$ as in Figure 2.2. Find an inverse function $x_2 = f_1^{-1}(x_1)$ and construct an auxiliary function $g(x_1) = f_1^{-1}(x_1) - f_2(x_1)$ that measures the distance between two best responses. It remains to show that $g(x_1)$ crosses zero only once since this would directly imply a single crossing point of $f_1(x_1)$ and $f_2(x_2)$. Suppose we could show that every time $g(x_1)$ crosses zero, it does so *from below*. If that is the case, we are assured there is only a single crossing: it is impossible for a continuous function to cross zero more than once from below because it would also have to cross zero from above somewhere. It can be shown that the function $g(x_1)$ crosses zero only from below if the slope of $g(x_1)$ at the crossing point is positive as follows

$$\frac{\partial g(x_1)}{\partial x_1} = \frac{\partial f_1^{-1}(x_1)}{\partial x_1} - \frac{\partial f_2(x_1)}{\partial x_1} = \frac{1}{\frac{\partial f_2(x_2)}{\partial x_2}} - \frac{\partial f_2(x_1)}{\partial x_1} > 0,$$

which holds if (2.5) holds. Hence, in a two-player game condition (2.5) is sufficient for the uniqueness of the NE. Note that condition (2.5) trivially holds in the newsvendor game since each slope is less than one and hence the multiplication of slopes is less than one as well *everywhere*. Index theory has been used by Netessine and Rudi (2001b) to show uniqueness of the NE in a retailer-wholesaler game when both parties stock inventory and sell directly to consumers and by Cachon and Kok (2002) and Cachon and Zipkin (1999).

2.5 Multiple equilibria

Many games are just not blessed with a unique equilibrium. The next best situation is to have a few equilibria. The worst situation is either to have an infinite number of equilibria or no equilibrium at all. The obvious problem with multiple equilibria is that the players may not know which equilibrium will prevail. Hence, it is entirely possible that

a non-equilibrium outcome results because one player plays one equilibrium strategy while a second player chooses a strategy associated with another equilibrium. However, if a game is repeated, then it is possible that the players eventually find themselves in one particular equilibrium. Furthermore, that equilibrium may not be the most desirable one.

If one does not want to acknowledge the possibility of multiple outcomes due to multiple equilibria, one could argue that one equilibrium is more reasonable than the others. For example, there may exist only one symmetric equilibrium and one may be willing to argue that a symmetric equilibrium is more focal than an asymmetric equilibrium. (See Mahajan and van Ryzin 1999a for an example). In addition, it is generally not too difficult to demonstrate the uniqueness of a symmetric equilibrium. If the players have unidimensional strategies, then the system of n first-order conditions reduces to a single equation and one need only show that there is a unique solution to that equation to prove the symmetric equilibrium is unique. If the players have m -dimensional strategies, $m > 1$, then finding a symmetric equilibrium reduces to determining whether a system of m equations has a unique solution (easier than the original system, but still challenging).

An alternative method to rule out some equilibria is to focus only on the Pareto optimal equilibrium, of which there may be only one. For example, in supermodular games the equilibria are Pareto rankable under an additional condition that each players' objective function is increasing in other players' strategies, i.e., there is a most preferred equilibrium by every player and a least preferred equilibrium by every player. (See Wang and Gerchak 2002 for an example). However, experimental evidence exists that suggests players do not necessarily gravitate to the Pareto optimal equilibrium as is demonstrated by Cachon and Camerer, 1996. Hence, caution is warranted with this argument.

2.6 Comparative statics in games

In GT models, just as in the non-competitive SCM models, many of the managerial insights and results are obtained through comparative statics such as monotonicity of the optimal decisions w.r.t. some parameter of the game.

2.6.1 The Implicit Functions Theorem approach. This approach works for both GT and single decision-maker applications as will become evident from the statement of the next theorem.

THEOREM 9. Consider the system of equations

$$\frac{\partial \pi_i(x_1, \dots, x_n, a)}{\partial x_i} = 0, \quad i = 1, \dots, n,$$

defining x_1^*, \dots, x_n^* as implicit functions of parameter a . If all derivatives are continuous functions and the Hessian (2.3) evaluated at x_1^*, \dots, x_n^* is non-zero, then the function $x^*(a)$, $R^1 \rightarrow R^n$ is continuous on a ball around x^* and its derivatives are found as follows:

$$\begin{pmatrix} \frac{\partial x_1^*}{\partial a} \\ \frac{\partial x_2^*}{\partial a} \\ \dots \\ \frac{\partial x_n^*}{\partial a} \end{pmatrix} = - \begin{pmatrix} \frac{\partial^2 \pi_1}{\partial x_1^2} & \frac{\partial^2 \pi_1}{\partial x_1 \partial x_2} & \dots & \frac{\partial^2 \pi_1}{\partial x_1 \partial x_n} \\ \frac{\partial^2 \pi_2}{\partial x_2 \partial x_1} & \frac{\partial^2 \pi_2}{\partial x_2^2} & \dots & \frac{\partial^2 \pi_2}{\partial x_2 \partial x_n} \\ \dots & \dots & \dots & \dots \\ \frac{\partial^2 \pi_n}{\partial x_n \partial x_1} & \frac{\partial^2 \pi_n}{\partial x_n \partial x_2} & \dots & \frac{\partial^2 \pi_n}{\partial x_n^2} \end{pmatrix}^{-1} \begin{pmatrix} \frac{\partial \pi_1}{\partial x_1 \partial a} \\ \frac{\partial \pi_1}{\partial x_2 \partial a} \\ \dots \\ \frac{\partial \pi_1}{\partial x_n \partial a} \end{pmatrix}. \quad (2.6)$$

Since the IFT is covered in detail in many non-linear programming books and its application to the GT problems is essentially the same, we do not delve further into this matter. In many practical problems, if $|H| \neq 0$ then it is instrumental to multiply both sides of the expression (2.6) by H^{-1} . That is justified because the Hessian is assumed to have a non-zero determinant to avoid the cumbersome task of inverting the matrix. The resulting expression is a system of n linear equations which have a closed form solution. See Netessine and Rudi (2001b) for such an application of the IFT in a two-player game and Bernstein and Federgruen (2000) in n -player games.

The solution to (2.6) in the case of two players is

$$\frac{\partial x_1^*}{\partial a} = - \frac{\frac{\partial^2 \pi_1}{\partial x_1 \partial a} \frac{\partial^2 \pi_2}{\partial x_2^2} - \frac{\partial^2 \pi_1}{\partial x_1 \partial x_2} \frac{\partial^2 \pi_2}{\partial x_2 \partial a}}{|H|}, \quad (2.7)$$

$$\frac{\partial x_2^*}{\partial a} = - \frac{\frac{\partial^2 \pi_1}{\partial x_1^2} \frac{\partial^2 \pi_2}{\partial x_2 \partial a} - \frac{\partial^2 \pi_1}{\partial x_1 \partial a} \frac{\partial^2 \pi_2}{\partial x_2 \partial x_1}}{|H|}. \quad (2.8)$$

Using our newsvendor game as an example, suppose we would like to analyze sensitivity of the equilibrium solution to changes in r_1 so let $a = r_1$. Notice that $\partial \pi_2 / \partial Q_2 \partial r_1 = 0$ and also that the determinant of the Hessian is positive. Both expressions in the numerator of (2.7) are positive as well so that $\partial Q_1^* / \partial r_1 > 0$. Further, the numerator of (2.8) is negative so that $\partial Q_2^* / \partial r_1 < 0$. Both results are intuitive.

Solving a system of n equations analytically is generally cumbersome and one may have to use Kramer's rule or analyze an inverse of H

instead, see Bernstein and Federgruen (2000) for an example. The only way to avoid this complication is to employ supermodular games as described below. However, the IFT method has an advantage that is not enjoyed by supermodular games: it can handle constraints of any form. That is, any constraint on the players' strategy spaces of the form $g_i(x_i) \leq 0$ or $g_i(x_i) = 0$ can be added to the objective function by forming a Lagrangian:

$$L_i(x_1, \dots, x_n, \lambda_i) = \pi_i(x_1, \dots, x_n) - \lambda_i g_i(x_i).$$

All analysis can then be carried through the same way as before with the only addition being that the Lagrange multiplier λ_i becomes a decision variable. For example, let's assume in the newsvendor game that the two competing firms stock inventory at a warehouse. Further, the amount of space available to each company is a function of the total warehouse capacity C , e.g., $g_i(Q_i) \leq C$. We can construct a new game where each retailer solves the following problem:

$$\max_{Q_i \in \{g_i(Q_i) \leq C\}} E_D \left[r_i \min \left(D_i + (D_j - Q_j)^+, Q_i \right) - c_i Q_i \right], \quad i = 1, 2.$$

Introduce two Lagrange multipliers, λ_i , $i = 1, 2$ and re-write the objective functions as

$$\max_{Q_i, \lambda_i} L(Q_i, \lambda_i, Q_j) = E_D \left[r_i \min \left(D_i + (D_j - Q_j)^+, Q_i \right) - c_i Q_i - \lambda_i (g_i(Q_i) - C) \right].$$

The resulting four optimality conditions can be analyzed using the IFT the same way as has been demonstrated previously.

2.6.2 Supermodular games approach. In some situations, supermodular games provide a more convenient tool for comparative statics.

THEOREM 11. *Consider a collection of supermodular games on R^n parameterized by a parameter a . Further, suppose $\partial^2 \pi_i / \partial x_i \partial a \geq 0$ for all i . Then the largest and the smallest equilibria are increasing in a .*

Roughly speaking, a sufficient condition for monotone comparative statics is supermodularity of players' payoffs in strategies *and* a parameter. Note that, if there are multiple equilibria, we cannot claim that every equilibrium is monotone in a ; rather, a set of all equilibria is monotone in the sense of Theorem 10. A convenient way to think about

the last Theorem is through the augmented Hessian:

$$\begin{vmatrix} \frac{\partial^2 \pi_1}{\partial x_1^2} & \frac{\partial^2 \pi_1}{\partial x_1 \partial x_2} & \cdots & \frac{\partial^2 \pi_1}{\partial x_1 \partial x_n} & \frac{\partial^2 \pi_1}{\partial x_1 \partial a} \\ \frac{\partial^2 \pi_2}{\partial x_2 \partial x_1} & \frac{\partial^2 \pi_2}{\partial x_2^2} & \cdots & \frac{\partial^2 \pi_2}{\partial x_2 \partial x_n} & \frac{\partial^2 \pi_2}{\partial x_2 \partial a} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \frac{\partial^2 \pi_n}{\partial x_n \partial x_1} & \frac{\partial^2 \pi_n}{\partial x_n \partial x_2} & \cdots & \frac{\partial^2 \pi_n}{\partial x_n^2} & \frac{\partial^2 \pi_n}{\partial x_n \partial a} \\ \frac{\partial^2 \pi_1}{\partial x_1 \partial a} & \frac{\partial^2 \pi_1}{\partial x_2 \partial a} & \cdots & \frac{\partial^2 \pi_1}{\partial x_n \partial a} & \frac{\partial^2 \pi_1}{\partial a^2} \end{vmatrix}.$$

Roughly, if all *off-diagonal* elements of this matrix are positive, then the monotonicity result holds (signs of diagonal elements do not matter and hence concavity is not required). To apply this result to competing newsvendors we will analyze sensitivity of equilibrium inventories (Q_i^*, Q_j^*) to r_i . First, transform the game to strategies (Q_i, y) so that the game is supermodular and find cross-partial derivatives

$$\begin{aligned} \frac{\partial^2 \pi_i}{\partial Q_i \partial r_i} &= \Pr(D_i + (D_j - Q_j)^+ > Q_i) \geq 0, \\ \frac{\partial \pi_j}{\partial y \partial r_i} &= 0 \geq 0, \end{aligned}$$

so that (Q_i^*, y^*) are both increasing in r_i or Q_i^* is increasing and Q_j^* is decreasing in r_i just as we have already established using the IFT.

The simplicity of the argument (once supermodular games are defined) as compared to the machinery required to derive the same result using the IFT is striking. Such simplicity has attracted much attention in SCM and has resulted in extensive applications of supermodular games. Examples include Cachon (2001), Corbett and DeCroix (2001), Netessine and Shumsky (2001) and Netessine and Rudi (2001b), to name just a few. There is, however, an important limitation to the use of Theorem 10: it cannot handle many constraints as IFT can. Namely, the decision space must be a lattice to apply supermodularity i.e., it must include its coordinate-wise maximum and minimum. Hence, a constraint of the form $x_i \leq b$ can be handled but a constraint $x_i + x_j \leq b$ cannot since points $(x_i, x_j) = (b, 0)$ and $(x_i, x_j) = (0, b)$ are within the constraint but the coordinate-wise maximum of these two points (b, b) is not. Notice that to avoid dealing with this issue in detail we stated in the theorems that the strategy space should all be R^n . Since in many SCM applications there are constraints on the players' strategies, supermodularity must be applied with care.

3. Dynamic games

While many SCM models are static, including all newsvendor-based models, a significant portion of the SCM literature is devoted to dynamic models in which decisions are made over time. In most cases the solution concept for these games is similar to the backwards induction used when solving dynamic programming problems. There are, however, important differences as will be clear from the discussion of repeated games. As with dynamic programming problems, we continue to focus on the games of complete information, i.e., at each move in the game all players know the full history of play.

3.1 Sequential moves: Stackelberg equilibrium concept

The simplest possible dynamic game was introduced by Stackelberg (1934). In a Stackelberg duopoly model, player 1, the Stackelberg leader, chooses a strategy first and then player 2, the Stackelberg follower, observes this decision and makes his own strategy choice. Since in many SCM models the upstream firm, e.g., the wholesaler, possesses certain power over the typically smaller downstream firm, e.g., the retailer, the Stackelberg equilibrium concept has found many applications in SCM literature. We do not address the issues of who should be the leader and who should be the follower, leaving those issues to Chapter 11.

To find an equilibrium of a Stackelberg game which is often called the Stackelberg equilibrium we need to solve a dynamic multi-period problem via backwards induction. We will focus on a two-period problem for analytical convenience. First, find the solution $x_2^*(x_1)$ for the second player as a response to any decision made by the first player:

$$x_2^*(x_1) : \frac{\partial \pi_2(x_2, x_1)}{\partial x_2} = 0.$$

Next, find the solution for the first player anticipating the response by the second player:

$$\frac{d\pi_1(x_1, x_2^*(x_1))}{dx_1} = \frac{\partial \pi_1(x_1, x_2^*)}{\partial x_1} + \frac{\partial \pi_1(x_1, x_2)}{\partial x_2} \frac{\partial x_2^*}{\partial x_1} = 0.$$

Intuitively, the first player chooses the best possible point on the second player's best response function. Clearly, the first player can choose a NE, so the leader is always at least as well off as he would be in NE. Hence, if a player were allowed to choose between making moves simultaneously or being a leader in a game with complete information he would

always prefer to be the leader. However, if new information is revealed after the leader makes a play, then it is not always advantageous to be the leader.

Whether the follower is better off in the Stackelberg or simultaneous move game depends on the specific problem setting. See Netessine and Rudi (2001a) for examples of both situations and comparative analysis of Stackelberg vs NE; see also Wang and Gerchak (2002) for a comparison between the leader vs follower roles in a decentralized assembly model. For example, consider the newsvendor game with sequential moves. The best response function for the second player remains the same as in the simultaneous move game:

$$Q_2^*(Q_1) = F_{D_2+(D_1-Q_1)^+}^{-1} \left(\frac{r_2 - c_2}{r_2} \right).$$

For the leader the optimality condition is

$$\begin{aligned} \frac{d\pi_1(Q_1, Q_2^*(Q_1))}{dQ_1} &= r_1 \Pr(D_1 + (D_2 - Q_2)^+ > Q_1) - c_1 \\ &\quad - r_1 \Pr(D_1 + (D_2 - Q_2)^+ < Q_1, D_2 > Q_2) \frac{\partial Q_2^*}{\partial Q_1} \\ &= 0, \end{aligned}$$

where $\partial Q_2^*/\partial Q_1$ is the slope of the best response function found in (2.1). Existence of a Stackelberg equilibrium is easy to demonstrate given the continuous payoff functions. However, uniqueness may be considerably harder to demonstrate. A sufficient condition is quasi-concavity of the leader's profit function, $\pi_1(x_1, x_2^*(x_1))$. In the newsvendor game example, this implies the necessity of finding derivatives of the density function of the demand distribution as is typical for many problems involving uncertainty. In stochastic models this is feasible with certain restrictions on the demand distribution. See Lariviere and Porteus (2001) for an example with a supplier that establishes the wholesale price and a newsvendor that then chooses an order quantity and Cachon (2003) for the reverse scenario in which a retailer sets the wholesale price and buys from a newsvendor supplier. See Netessine and Rudi (2001a) for a Stackelberg game with a wholesaler choosing a stocking quantity and the retailer deciding on promotional effort. One can further extend the Stackelberg equilibrium concept into multiple periods, see Erhun et al. (2000) and Anand et al. (2002) for examples.

3.2 Simultaneous moves: repeated and stochastic games

A different type of dynamic game arises when both players take actions in multiple periods. Since inventory models used in SCM literature often involve inventory replenishment decisions that are made over and over again, multi-period games should be a logical extension of these inventory models. Two major types of multiple-period games exist: without and with time dependence.

In the multi-period game without time dependence the exact same game is played over and over again hence the term *repeated* games. The strategy for each player is now a sequence of actions taken in all periods. Consider one repeated game version of the competing newsvendor game in which the newsvendor chooses a stocking quantity at the start of each period, demand is realized and then leftover inventory is salvaged. In this case, there are no links between successive periods other than the players' memory about actions taken in all the previous periods. Although repeated games have been extensively analyzed in economics literature, it is awkward in a SCM setting to assume that nothing links successive games; typically in SCM there is some transfer of inventory and/or backorders between periods. As a result, repeated games thus far have not found many applications in the SCM literature. Exceptions are Debo (1999), Taylor and Plambeck (2003) and Ren et al. (2003) in which reputational effects are explored as means of supply chain coordination in place of the formal contracts.

A fascinating feature of repeated games is that the set of equilibria is much larger than the set of equilibria in a static game and may include equilibria that are not possible in the static game. At first, one may assume that the equilibrium of the repeated game would be to play the same static NE strategy in each period. This is, indeed, an equilibrium but only one of many. Since in repeated games the players are able to condition their behavior on the observed actions in the previous periods, they may employ so-called *trigger strategies*: the player will choose one strategy until the opponent changes his play at which point the first player will change the strategy. This *threat* of reverting to a different strategy may even induce players to achieve the best possible outcome, i.e., the centralized solution, which is called an *implicit collusion*. Many such threats are, however, non-credible in the sense that once a part of the game has been played, such a strategy is not an equilibrium anymore for the remainder of the game, as is the case in our example in Figure

2.1. To separate out credible threats from non-credible, Selten (1965) introduced the notion of a *subgame-perfect equilibrium*. See Hall and Porteus (2000) and van Mieghem and Dada (1999) for solutions involving subgame-perfect equilibria in dynamic games.

Subgame-perfect equilibria reduce the equilibrium set somewhat. However, infinitely-repeated games are still particularly troublesome in terms of multiplicity of equilibria. The famous Folk theorem⁶ proves that any convex combination of the feasible payoffs is attainable in the infinitely repeated game as an equilibrium, implying that “virtually anything” is an equilibrium outcome⁷. See Debo (1999) for the analysis of a repeated game between the wholesaler setting the wholesale price and the newsvendor setting the stocking quantity.

In time-dependent multi-period games players’ payoffs in each period depend on the actions in the previous as well as current periods. Typically the payoff structure does not change from period to period (so called stationary payoffs). Clearly, such setup closely resembles multi-period inventory models in which time periods are connected through the transfer of inventories and backlogs. Due to this similarity, time-dependent games have found applications in SCM literature. We will only discuss one type of time-dependent multi-period games, *stochastic games* or *Markov games*, due to their wide applicability in SCM. See also Majumder and Groenevelt (2001b) for the analysis of deterministic time-dependent multi-period games in reverse logistics supply chains. Stochastic games were developed by Shapley (1953a) and later by Sobel (1971), Kirman and Sobel (1974) and Heyman and Sobel (1984). The theory of stochastic games is also extensively covered in Filar and Vrieze (1996).

The setup of the stochastic game is essentially a combination of a static game and a Markov Decisions Process: in addition to the set of players with strategies which is now a vector of strategies, one for each period, and payoffs, we have a set of states and a transition mechanism $p(s'|s, x)$, probability that we transition from state s to state s' given action x . Transition probabilities are typically defined through random demand occurring in each period. The difficulties inherent in considering non-stationary inventory models are passed over to the game-theoretic extensions of these models, so a standard simplifying assumption is that demands are independent and identical across periods. When only a single decision-maker is involved, such an assumption leads to a unique stationary solution (e.g., stationary inventory policy of some form: order-

up-to, S-s, etc.). In a GT setting, however, things get more complicated; just as in the repeated games described above, non-stationary equilibria, e.g., trigger strategies, are possible. A standard approach is to consider just one class of equilibria – e.g., stationary – since non-stationary policies are hard to implement in practice and they are not always intuitively appealing. Hence, with the assumption that the policy is stationary the stochastic game reduces to an equivalent static game and equilibrium is found as a sequence of NE in an appropriately modified single-period game. Another approach is to focus on “Markov” or “state-space” strategies in which the past influences the future through the state variables but not through the history of the play. A related equilibrium concept is that of *Markov Perfect Equilibrium* (MPE) which is simply a profile of Markov strategies that yields a Nash equilibrium in every subgame. The concept of MPE is discussed in Chapter 13 of Fudenberg and Tirole (1991). See also Tayur and Yang (2002) for the application of this concept.

To illustrate, consider an infinite-horizon variant of the newsvendor game with lost sales in each period and inventory carry-over to the subsequent period, see Netessine et al. 2002 for complete analysis. The solution to this problem in a non-competitive setting is an order-up-to policy. In addition to unit-revenue r and unit-cost c we introduce inventory holding cost h incurred by a unit carried over to the next period and a discount factor β . Also denote by x_i^t the inventory position at the beginning of the period and by y_i^t the order-up-to quantity. Then the infinite-horizon profit of each player is

$$\pi_i(x^1) = E \sum_{t=1}^{\infty} \beta_i^{t-1} \left[r_i \min \left(y_i^t, D_i^t + (D_j^t - y_j^t)^+ \right) - h_i \left(y_i^t - D_i^t - (D_j^t - y_j^t)^+ \right)^+ - c_i Q_i^t \right],$$

with the inventory transition equation

$$x_i^{t+1} = \left(y_i^t - D_i^t - (D_j^t - y_j^t)^+ \right)^+.$$

Using the standard manipulations from Heyman and Sobel (1984), this objective function can be converted to

$$\pi_i(x^1) = c_i x_i^1 + \sum_{t=1}^{\infty} \beta_i^{t-1} G_i^t(y_i^t), \quad i = 1, 2,$$

where $G_i^t(y_i^t)$ is a single-period objective function

$$\begin{aligned} G_i^t(y_i^t) &= E \left[(r_i - c_i) \left(D_i^t + (D_j^t - y_j^t)^+ \right) \right. \\ &\quad - (r_i - c_i) \left(D_i^t + (D_j^t - y_j^t)^+ - y_i^t \right)^+ \\ &\quad \left. - (h_i + c_i(1 - \beta_i)) \left(y_i^t - D_i^t - (D_j^t - y_j^t)^+ \right)^+ \right], \quad i = 1, 2, t = 1, 2, \dots \end{aligned}$$

Assuming demand is stationary and independently distributed across periods $D_i = D_i^t$ we further obtain that $G_i^t(y_i^t) = G_i(y_i^t)$ since the single-period game is the same in each period. By restricting consideration to the stationary inventory policy $y_i = y_i^t$, $t = 1, 2, \dots$, we can find the solution to the multi-period game as a sequence of the solutions to a single-period game $G_i(y_i)$ which is

$$y_i^* = F_{D_i + (D_j - y_j^*)^+}^{-1} \left(\frac{r_i - c_i}{r_i + h_i - c_i \beta_i} \right), \quad i = 1, 2.$$

With the assumption that the equilibrium is stationary, one could argue that stochastic games are no different from static games; except for a small change in the right-hand side reflecting inventory carry-over and holding costs, the solution is essentially the same. However, more elaborate models capture some effects that are not present in static games but can be envisioned in stochastic games. For example, if we were to introduce backlogging in the above model, a couple of interesting situations would arise: a customer may backlog the product with either the first or with the second competitor he visits if both are out of stock. These options introduce the behavior that is observed in practice but cannot be modeled within the static game (see Netessine et al. 2002 for detailed analysis) since firms' inventory decisions affect their demand in the future. Among other applications of stochastic games are papers by Cachon and Zipkin (1999) analyzing a two-echelon game with the wholesaler and the retailer making stocking decisions, Bernstein and Federgruen (2002) analyzing price and service competition, Netessine and Rudi (2001a) analyzing the game with the retailer exerting sales effort and the wholesaler stocking the inventory and van Mieghem and Dada (1999) studying a two-period game with capacity choice in the first period and production decision under the capacity constraint in the second period.

3.3 Differential games

So far we have described dynamic games in discrete time, i.e., games involving a sequence of decisions that are separated in time. *Differential games* provide a natural extension for decisions that have to be made continuously. Since many SC models rely on continuous-time processes, it is natural to assume that differential games should find a variety of applications in SCM literature. However, most SCM models include stochasticity in one form or another. At the same time, due to the mathematical difficulties inherent in differential games, we are only aware of deterministic differential GT models in SCM. Although theory for stochastic differential games does exist, applications are quite limited, see Basar and Olsder (1995). Marketing and economics have been far more successful in applying differential games since deterministic models are standard in these areas. Hence, we will only briefly outline some new concepts necessary to understand the theory of differential games.

The following is a simple example of a differential game taken from Kamien and Schwartz (2000). Suppose two players indexed by $i = 1, 2$ are engaged into production and sales of the same product. Firms choose production levels $u_i(t)$ at any moment of time and incur total cost $C_i(u_i) = cu_i + u_i^2/2$. The price in the market is determined as per Cournot competition. Typically, this would mean that $p(t) = a - u_1(t) - u_2(t)$. However, the twist in this problem is that if the production level is changed, price adjustments are not instantaneous. Namely, there is a parameter s , referred to as the speed of price adjustment, so that the price is adjusted according to the following differential equation:

$$p'(t) = s [a - u_1(t) - u_2(t) - p(t)], p(0) = p_0.$$

Finally, each firm maximizes

$$\pi_i = \int_0^{\infty} e^{-rt} (p(t)u_i(t) - C_i(u_i(t))) dt, i = 1, 2.$$

The standard tools needed to analyze differential games are the calculus of variations or optimal control theory, see Kamien and Schwartz (2000). In a standard optimal control problem a single decision-maker sets the control variable that affects the state of the system. In contrast, in differential games several players select control variables that may affect a common state variable and/or payoffs of all players. Hence, differential games can be looked at as a natural extension of the optimal control theory. In this section we will consider two distinct types of player strategies: *open-loop* and *closed-loop* which is also sometimes

called *feedback*. In the open-loop strategy the players select their decisions or control variables once at the beginning of the game and do not change them so that the control variables are only functions of time and do not depend on the other players' strategies. Open-loop strategies are simpler in that they can be found through the straightforward application of optimal control which makes them quite popular. Unfortunately, an open-loop strategy may not be subgame-perfect. On the contrary, in a closed-loop strategy the player bases his strategy on current time and the states of both players' systems. Hence, feedback strategies are subgame-perfect: if the game is stopped at any time, for the remainder of the game the same feedback strategy will be optimal, which is consistent with the solution to the dynamic programming problems that we employed in the stochastic games section. The concept of a feedback strategy is more satisfying, but is also more difficult to analyze. In general, optimal open-loop and feedback strategies differ, but they may coincide in some games.

Since it is hard to apply differential game theory in stochastic problems, we cannot utilize the competitive newsvendor problem to illustrate the analysis. Moreover, the analysis of even the most trivial differential game is somewhat involved mathematically so we will limit our survey to stating and contrasting optimality conditions in the cases of open-loop and closed-loop NE. Stackelberg equilibrium models do exist in differential games as well but are rarer, see Basar and Olsder (1995). Due to mathematical complexity, games with more than two players are hardly ever analyzed. In a differential game with two players, each player is endowed with a control $u_i(t)$ that the player uses to maximize the objective function π_i

$$\max_{u_i(t)} \pi_i(u_i, u_j) = \max_{u_i(t)} \int_0^T f_i(t, x_i(t), x_j(t), u_i(t), u_j(t)) dt,$$

where $x_i(t)$ is a state variable describing the state of the system. The state of the system evolves according to the differential equation

$$x_i'(t) = g_i(t, x_i(t), x_j(t), u_i(t), u_j(t)),$$

which is the analog of the inventory transition equation in the multi-period newsvendor problem. Finally, there are initial conditions $x_i(0) = x_{i0}$.

The open-loop strategy implies that each players' control is only a function of time, $u_i = u_i(t)$. A feedback strategy implies that each players' control is also a function of state variables, $u_i = u_i(t, x_i(t), x_j(t))$.

As in the static games, NE is obtained as a fixed point of the best response mapping by simultaneously solving a system of first-order optimality conditions for the players. Recall that to find the optimal control we first need to form a Hamiltonian. If we were to solve two individual non-competitive optimization problems, the Hamiltonians would be $H_i = f_i + \lambda_i g_i$, $i = 1, 2$, where $\lambda_i(t)$ is an adjoint multiplier. However, with two players we also have to account for the state variable of the opponent so that the Hamiltonian becomes

$$H_i = f_i + \lambda_i^1 g_i + \lambda_i^2 g_j, \quad i, j = 1, 2.$$

To obtain the necessary conditions for the open-loop NE we simply use the standard necessary conditions for any optimal control problem:

$$\frac{\partial H_1}{\partial u_1} = 0, \quad \frac{\partial H_2}{\partial u_2} = 0, \quad (2.9)$$

$$\frac{\partial \lambda_1^1}{\partial t} = -\frac{\partial H_1}{\partial x_1}, \quad \frac{\partial \lambda_1^2}{\partial t} = -\frac{\partial H_1}{\partial x_2}, \quad (2.10)$$

$$\frac{\partial \lambda_2^1}{\partial t} = -\frac{\partial H_2}{\partial x_2}, \quad \frac{\partial \lambda_2^2}{\partial t} = -\frac{\partial H_2}{\partial x_1}. \quad (2.11)$$

For the feedback equilibrium the Hamiltonian is the same as for the open-loop strategy. However, the necessary conditions are somewhat different:

$$\frac{\partial H_1}{\partial u_1} = 0, \quad \frac{\partial H_2}{\partial u_2} = 0, \quad (2.12)$$

$$\frac{\partial \lambda_1^1}{\partial t} = -\frac{\partial H_1}{\partial x_1} - \frac{\partial H_1}{\partial u_2} \frac{\partial u_2^*}{\partial x_1}, \quad \frac{\partial \lambda_1^2}{\partial t} = -\frac{\partial H_1}{\partial x_2} - \frac{\partial H_1}{\partial u_2} \frac{\partial u_2^*}{\partial x_2}, \quad (2.13)$$

$$\frac{\partial \lambda_2^1}{\partial t} = -\frac{\partial H_2}{\partial x_2} - \frac{\partial H_2}{\partial u_1} \frac{\partial u_1^*}{\partial x_2}, \quad \frac{\partial \lambda_2^2}{\partial t} = -\frac{\partial H_2}{\partial x_1} - \frac{\partial H_2}{\partial u_1} \frac{\partial u_1^*}{\partial x_1}. \quad (2.14)$$

Notice that the difference is captured by an extra term on the right when we compare (2.10) and (2.13) or (2.11) and (2.14). The difference is due to the fact that the optimal control of each player under the feedback strategy depends on $x_i(t)$, $i = 1, 2$. Hence, when differentiating the Hamiltonian to obtain equations (2.13) and (2.14) we have to account for such dependence (note also that two terms disappear when we use (2.12) to simplify).

As we mentioned earlier, there are numerous applications of differential games in economics and marketing, especially in the area of dynamic pricing, see, e.g., Eliashberg and Jeuland (1986). Eliashberg and Steinberg (1987), Desai (1992) and Desai (1996) use the open-loop Stackelberg

equilibrium concept in a marketing-production game with the manufacturer and the distributor. Gaimon (1989) uses both open and closed-loop NE concepts in a game with two competing firms choosing prices and production capacity when the new technology reduces firms' costs. Mukhopadhyay and Kouvelis (1997) consider a duopoly with firms competing on prices and quality of design and derive open and closed-loop NE.

4. Cooperative games

The subject of cooperative games first appeared in the seminal work of von Neumann and Morgenstern (1944). However, for a long time cooperative game theory did not enjoy as much attention in the economics literature as non-cooperative GT. Papers employing cooperative GT to study SCM had been scarce, but are becoming more popular. This trend is probably due to the prevalence of bargaining and negotiations in SC relationships.

Cooperative GT involves a major shift in paradigms as compared to non-cooperative GT: the former focuses on the outcome of the game in terms of the value created through cooperation of a subset of players but does not specify the actions that each player will take, while the latter is more concerned with the specific actions of the players. Hence, cooperative GT allows us to model outcomes of complex business processes that otherwise might be too difficult to describe, e.g., negotiations, and answers more general questions, e.g., how well is the firm positioned against competition (see Brandenburger and Stuart 1996). However, there are also limitations to cooperative GT, as we will later discuss.

In what follows, we will cover transferable utility cooperative games (players can share utility via side payments) and two solution concepts: the core of the game and the Shapley value, and also biform games that have found several applications in SCM. Not covered are alternative concepts of value, e.g., nucleous and the σ -value, and games with non-transferable utility that have not yet found application in SCM. Material in this section is based mainly on Stuart (2001) and Moulin (1995). Perhaps the first paper employing cooperative games in SCM is Wang and Parlar (1994) who analyze the newsvendor game with three players, first in a non-cooperative setting and then under cooperation with and without Transferable Utility.

4.1 Games in characteristic form and the core of the game

Recall that the non-cooperative game consists of a set of players with their strategies and payoff functions. In contrast, the cooperative game, which is also called the game in characteristic form, consists of the set of players N with subsets or *coalitions* $S \subseteq N$ and a *characteristic function* $v(S)$ that specifies a (maximum) value (which we assume is a real number) created by any subset of players in N , i.e., the total pie that members of a coalition can create and divide. The specific actions that players have to take to create this value are not specified: the characteristic function only defines the total value that can be created by utilizing all players' resources. Hence, players are free to form any coalitions that are beneficial to them and no player is endowed with power of any sort. Furthermore, the value a coalition creates is independent of the coalitions and actions taken by the non-coalition members. This decoupling of payoffs is natural in political settings (e.g., the majority gets to choose the legislation), but it far more problematic in competitive markets. For example, in the context of cooperative game theory, the value HP and Compaq can generate by merging is independent of the actions taken by Dell, Gateway, IBM, Ingram Micron, etc.⁸

A frequently used solution concept in cooperative GT is the *core of the game*:

DEFINITION 5: *The utility vector π_1, \dots, π_N is in the core of the cooperative game if $\forall S \subset N$, $\sum_{i \in S} \pi_i \geq v(S)$ and $\sum_{i \in N} \pi_i \geq v(N)$*

A utility vector is in the core if the total utility of every possible coalition is at least as large as the coalition's value, i.e., there does not exist a coalition of players that could make all of its members at least as well off and one member strictly better off.

As is true for NE, the core of the game may not exist, i.e., it may be empty, and the core is often not unique. Existence of the core is an important issue because with an empty core it is difficult to predict what coalitions would form and what value each player would receive. If the core exists, then the core typically specifies a range of utilities that a player can appropriate, i.e., competition alone does not fully determine the players' payoffs. What utility each player will actually receives is undetermined: it may depend on the details of the residual bargaining process, a source of criticism of the core concept. (Biform games, de-

scribed below, provide one possible resolution of this indeterminacy.)

In terms of specific applications to the SCM, Hartman et al. (2000) considered the newsvendor centralization game, i.e., a game in which multiple retailers decide to centralize their inventory and split profits resulting from the benefits of risk pooling. Hartman et al. (2000) further show that this game has a non-empty core under certain restrictions on the demand distribution. Muller et al. (2002) relax these restrictions and show that the core is always non-empty. Further, Muller et al. (2002) give a condition under which the core is a singleton.

4.2 Shapley value

The concept of the core, though intuitively appealing, also possesses some unsatisfying properties. As we mentioned, the core might be empty or indeterministic⁹. As it is desirable to have a unique NE in non-cooperative games, it is desirable to have a solution concept for cooperative games that results in a unique outcome. Shapley (1953b) offered an axiomatic approach to a solution concept that is based on three axioms. First, the value of a player should not change due to permutations of players, i.e., only the role of the player matters and not names or indices assigned to players. Second, if a player's added value to the coalition is zero then this player should not get any profit from the coalition, or, in other words, only players generating added value should share the benefits. (A player's added value is the difference between the coalition's value with that player and without that player.) Those axioms are intuitive, but the third is far less so. The third axiom requires additivity of payoffs: if v_1 and v_2 are characteristic functions in any two games, and if q_1 and q_2 are a player's Shapley value in these two games, then the player's Shapley value in the composite game, $v_1 + v_2$, must be $q_1 + q_2$. This is not intuitive because it is not clear by what is meant by a composite game. Nevertheless, Shapley (1953b) demonstrates that there is a unique value for each player, called the Shapley value, that satisfies all three axioms.

THEOREM 11. *The Shapley value, π_i , for player i in an N -person non-cooperative game with transferable utility is :*

$$\pi_i = \sum_{S \subseteq N \setminus i} \frac{|S|! (|N| - |S| - 1)!}{|N|!} (v(S \cup \{i\}) - v(S)).$$

The Shapley value assigns to each player his marginal contribution ($v(S \cup \{i\}) - v(S)$) when S is a random coalition of agents preceding i and the ordering is

drawn at random. To explain further, (see Myerson 1997), suppose players are picked randomly to enter into a coalition. There are $|N|!$ different orderings for all players, and for any set S that does not contain player i there are $|S|!(|N| - |S| - 1)!$ ways to order players so that all of the players in S are picked ahead of player i . If the orderings are equally likely, there is a probability of $|S|!(|N| - |S| - 1)!/|N|!$ that when player i is picked he will find S players in the coalition already. The marginal contribution of adding player i to coalition S is $(v(S \cup \{i\}) - v(S))$. Hence, the Shapley value is nothing more than a marginal expected contribution of adding player i to the coalition.

Because the Shapley value is unique, it has found numerous applications in economics and political sciences. So far, however, SCM applications are scarce: except for discussion in Granot and Sosis (2001) we are not aware of any other papers employing the concept of the Shapley value. Although uniqueness of the Shapely value is a convenient feature, caution should surely be taken with Shapely value: the Shapely value need not be in the core, hence, although the Shapely is appealing from the perspective of fairness, it may not be a reasonable prediction of the outcome of a game (i.e., because it is not in the core, there exists some subset of players that can deviate and improve their lots).

4.3 Biform games

From the SCM point of view, cooperative games are somewhat unsatisfactory in that they do not explicitly describe the equilibrium actions taken by the players that is often the key in SC models. *Biform games*, developed by Brandenburger and Stuart (2003), compensate to some extent for this shortcoming.

A biform game can be thought of as a non-cooperative game with cooperative games as outcomes and those cooperative games lead to specific payoffs. Similar to the non-cooperative game, the biform game has a set of players N , a set of strategies for each player and also a cost function associated with each strategy (cost function is optional – we include it since most SCM applications of biform games involve cost functions). The game begins by players making choices from among their strategies and incurring costs. After that, a cooperative game occurs in which the characteristic value function depends on the chosen actions. Hopefully the core of each possible cooperative game is non-empty, but it is also unlikely to be unique. As a result, there is not specific outcome of the cooperative sub-game, i.e., it is not immediately clear what value each player can expect. The proposed solution is that each player is assigned a confidence index, $\alpha_i \in [0, 1]$, and the α_i s are

common knowledge. Each player then expects to earn in each possible cooperative game a weighted average of the minimum and maximum values in the core, with α_i being the weight. For example, if $\alpha_i = 0$, then the player earns the minimum value in the core and if $\alpha_i = 1$ then the player earns the maximum value in the core. Once a specific value is assigned to each player for each cooperative sub-game, the first stage non-cooperative game can be analyzed just like any other non-cooperative game.

Biform games have been successfully adopted in several SCM papers. Anupindi et al. (2001) consider a game where multiple retailers stock at their own locations as well as at several centralized warehouses. In the first (non-cooperative) stage retailers make stocking decisions. In the second (cooperative) stage retailers observe demand and decide how much inventory to transship among locations to better match supply and demand and how to appropriate the resulting additional profits. Anupindi et al. (2001) conjecture that a characteristic form of this game has an empty core. However, the biform game has a non-empty core and they find the allocation of rents based on dual prices that is in the core. Moreover, they find an allocation mechanism that is in the core and that allows them to achieve coordination, i.e., the first-best solution. Granot and Sosic (2001) analyze a similar problem but allow retailers to hold back the residual inventory. In their model there are actually three stages: inventory procurement, decision about how much inventory to share with others and finally the transshipment stage. Plambeck and Taylor (2001a, 2001b) analyze two similar games between two firms that have an option of pooling their capacity and investments to maximize the total value. In the first stage, firms choose investment into effort that affects the market size. In the second stage, firms bargain over the division of the market and profits.

5. Signaling, Screening and Bayesian Games

So far we have considered only games in which the players are on “equal footing” with respect to information, i.e., each player knows every other player’s expected payoff with certainty for any set of chosen actions. However, such ubiquitous knowledge is rarely present in supply chains. One firm may have a better forecast of demand than another firm, or a firm may possess superior information regarding its own costs and operating procedures. Furthermore, a firm may know that another firm may have better information, and therefore choose actions that acknowledge this information shortcoming. Fortunately, game theory provides tools to study these rich issues, but, unfortunately, they do add

another layer of analytical complexity. This section briefly describes three types of games in which the information structure has a strategic role: signaling games, screening games and Bayesian games. Detailed methods for the analysis of these games are not provided. Instead, a general description is provided along with specific references to supply chain management papers that study these games.

5.1 Signaling Game

In its simplest form, a signaling game has two players, one of which has better information than the other and it is the player with the better information that makes the first move. For example, Cachon and Lariviere (2001) consider a model with one supplier and one manufacturer. The supplier must build capacity for a key component to the manufacturer's product, but the manufacturer has a better demand forecast than the supplier. In an ideal world the manufacturer would truthfully share her demand forecast with the supplier so that the supplier could build the appropriate amount of capacity. However, the manufacturer always benefits from a larger installed capacity in case demand turns out to be high but it is the supplier that bears the cost of that capacity. Hence, the manufacturer has an incentive to inflate her forecast to the supplier. The manufacturer's hope is that the supplier actually believes the rosy forecast and builds additional capacity. Unfortunately, the supplier is aware of this incentive to distort the forecast, and therefore should view the manufacturer's forecast with skepticism. The key issue is whether there is something the manufacturer should do to make her forecast convincing, i.e., credible.

While the reader should refer to Cachon and Lariviere (2001) for the details of the game, some definitions and concepts are needed to continue this discussion. The manufacturer's private information, or type, is her demand forecast. There is a set of possible types that the manufacturer could be and this set is known to the supplier, i.e., the supplier is aware of the possible forecasts, but is not aware of the manufacturer's actual forecast. Furthermore, at the start of the game the supplier and the manufacturer know the probability distribution over the set of types. We refer to this probability distribution as the supplier's belief regarding the types. The manufacturer chooses her action first which in this case is a contract offer and a forecast, the supplier updates his belief regarding the manufacturer's type given the observed action, and then the supplier chooses his action which in this case is the amount of capacity to build. If the supplier's belief regarding the manufacturer's

type is resolved to a single type after observing the manufacturer's action (i.e., the supplier assigns a 100% probability that the manufacturer is that type and a zero probability that the manufacturer is any other type) then the manufacturer has signaled a type to the supplier. The trick for the supplier is to ensure that the manufacturer has signaled her actual type.

While we are mainly interested in the set of contracts that credibly signal the manufacturer's type, it is worth beginning with the possibility that the manufacturer does not signal her type. In other words, the manufacturer chooses an action such that the action does not provide the supplier with additional information regarding the manufacturer's type. That outcome is called a *pooling equilibrium*, because the different manufacturer types behave in the same way, i.e., the different types are pooled into the same set of actions. As a result, Bayes' rule does not allow the supplier to refine his beliefs regarding the manufacturer's type.

A pooling equilibrium is not desirable from the perspective of supply chain efficiency because the manufacturer's type is not communicated to the supplier. Hence, the supplier does not choose the correct capacity given the manufacturer's actual demand forecast. However, this does not mean that both firms are disappointed with a pooling equilibrium. If the manufacturer's demand forecast is worse than average, then that manufacturer is quite happy with the pooling equilibrium because the supplier is likely to build more capacity than he would if he learned the manufacturer's true type. It is the manufacturer with a higher than average demand forecast that is disappointed with the pooling equilibrium because then the supplier is likely to underinvest in capacity.

A pooling equilibrium is often supported by the belief that every type will play the pooling equilibrium and any deviation from that play would only be done by a manufacturer with a low demand forecast. This belief can prevent the high demand manufacturer from deviating from the pooling equilibrium: a manufacturer with a high demand forecast would rather be treated as an average demand manufacturer (the pooling equilibrium) than a low demand manufacturer (if deviating from the pooling equilibrium). Hence, a pooling equilibrium can indeed be a NE in the sense that no player has a unilateral incentive to deviate given the strategies and beliefs chosen by the other players.

While a pooling equilibrium can meet the criteria of a NE, it nevertheless may not be satisfying. In particular, why should the supplier

believe that the manufacturer is a low type if the manufacturer deviates from the pooling equilibrium? Suppose the supplier were to believe a deviating manufacturer has a high demand forecast. If a high type manufacturer is better off deviating but a low type manufacturer is not better off, then only the high type manufacturer would choose such a deviation. The key part in this condition is that the low type is not better off deviating. In that case it is not reasonable for the supplier to believe the deviating manufacturer could only be a high type, so the supplier should adjust his belief. Furthermore, the high demand manufacturer should then deviate from the pooling equilibrium, i.e., this reasoning, which is called the intuitive criterion, breaks the pooling equilibrium, see Kreps (1990).

The contrast to a pooling equilibrium is a *separating* equilibrium which is also called a *signaling* equilibrium. With a separating equilibrium the different manufacturer types choose different actions, so the supplier is able to perfectly refine his belief regarding the manufacturer's type given the observed action. The key condition for a separating equilibrium is that only one manufacturer type is willing to choose the action designated for that type. If there is a continuum of manufacturer types, then it is quite challenging to obtain a separating equilibrium: it is difficult to separate two manufacturers that have nearly identical types. However, separating equilibria are more likely to exist if there is a finite number of discrete types.

There are two main issues with respect to separating equilibria: what actions lead to separating equilibrium and does the manufacturer incur a cost to signal, i.e., is the manufacturer's expected profit in the separating equilibrium lower than what it would be if the manufacturer's type were known to the supplier with certainty? In fact, these two issues are related: an ideal action for a high demand manufacturer is one that costlessly signals her high demand forecast. If a costless signal does not exist, then the goal is to seek the lowest cost signal.

Cachon and Lariviere (2001) demonstrate that whether a costless signal exists depends upon what commitments the manufacturer can impose on the supplier. For example, suppose the manufacturer dictates to the supplier a particular capacity level in the manufacturer's contract offer. Furthermore, suppose the supplier accepts that contract and by accepting the contract the supplier has essentially no choice but to build that level of capacity since the penalty for noncompliance is too severe. They refer to this regime as forced compliance. In that case there exist

many costless signals for the manufacturer. However, if the manufacturer's contract is not iron-clad, so the supplier could potentially deviate, which is referred to as voluntary compliance, then the manufacturer's signaling task becomes more complex.

One solution for a high demand manufacturer is to give a sufficiently large lump sum payment to the supplier: the high demand manufacturer's profit is higher than the low demand manufacturer's profit, so only a high demand manufacturer could offer that sum. This has been referred to as signaling by "burning money": only a firm with a lot of money can afford to burn that much money.

While burning money can work, it is not a smart signal: burning one unit of income hurts the high demand manufacturer as much as it hurts the low demand manufacturer. The signal works only because the high demand manufacturer has more units to burn. A better signal is a contract offer that is costless to a high demand manufacturer but expensive to a low demand manufacturer. A good example of such a signal is a minimum commitment. A minimum commitment is costly only if realized demand is lower than the commitment, because then the manufacturer is forced to purchase more units than desired. That cost is less likely for a high demand manufacturer, so in expectation a minimum commitment is costlier for a low demand manufacturer. Interestingly, Cachon and Lariviere (2001) show that a manufacturer would never offer a minimum commitment with perfect information, i.e., these contracts may be used in practice solely for the purpose of signaling information.

5.2 Screening

In a screening game the player that lacks information is the first to move. For example, in the screening game version of the supplier-manufacturer game described by Cachon and Lariviere (2001) the supplier makes the contract offer. In fact, the supplier offers a menu of contracts with the intention of getting the manufacturer to reveal her type via the contract selected in the menu. In the economics literature this is also referred to as *mechanism design*, because the supplier is in charge of designing a mechanism to learn the manufacturer's information. See Porteus and Whang (1999) for a screening game that closely resembles this one.

The space of potential contract menus is quite large, so large that it is not immediately obvious how to begin to find the supplier's opti-

mal menu. For example, how many contracts should be offered and what form should they take? Furthermore, for any given menu the supplier needs to infer for each manufacturer type which contract the type will choose. Fortunately, the *revelation principle* (Kreps 1990) provides some guidance.

The revelation principle begins with the presumption that a set of optimal mechanisms exists. Associated with each of these mechanisms is a NE that specifies which contract each manufacturer type chooses and the supplier's action given the chosen contract. With some of these equilibria it is possible that some manufacturer type chooses a contract that is not designated for that type. For example, the supplier intends the low demand manufacturer to choose one of the menu options, but instead the high demand manufacturer chooses that option. Even though this does not seem desirable, it is possible that this mechanism is still optimal in the sense that the supplier can do no better on average. The supplier ultimately cares only about expected profit, not the means by which that profit is achieved. Nevertheless, the revelation principle states that an optimal mechanism that involves deception (the wrong manufacturer chooses a contract) can be replaced by a mechanism that does not involve deception, i.e., there exists an equivalent mechanism that is truth-telling. Hence, in the hunt for an optimal mechanism it is sufficient to consider the set of revealing mechanisms: the menu of contracts is constructed such that each option is designated for a type and that type chooses that option.

Even though an optimal mechanism may exist for the supplier, this does not mean the supplier earns as much profit as he would if he knew the manufacturer's type. The gap between what a manufacturer earns with the menu of contracts and what the same manufacturer would earn if the supplier knew her type is called an information rent. A feature of these mechanisms is that separation of the manufacturer types goes hand in hand with a positive information rent, i.e., a manufacturer's private information allows the manufacturer to keep some rent that the manufacturer would not be able to keep if the supplier knew her type. Hence, even though there may be no cost to information revelation with a signaling game, the same is not true with a screening game.

There have been a number of applications of the revelation principle in the supply chain literature: e.g., Chen (2001) studies auction design in the context of supplier procurement contracts; Corbett (2001) studies

inventory contract design; Baiman et al. (2003) study procurement of quality in a supply chain.

5.3 Bayesian games

With a signaling game or a screening game actions occur sequentially so information can be revealed through the observation of actions. There also exist games with private information that do not involve signaling or screening. Consider the capacity allocation game studied by Cachon and Lariviere (1999). A single supplier has a finite amount of capacity. There are multiple retailers, and each knows his own demand but not the demand of the other retailers. The supplier announces an allocation rule, the retailers submit their orders and then the supplier produces and allocates units. If the retailers' total order is less than capacity, then each retailer receives his entire order. If the retailers' total order exceeds capacity, the supplier's allocation rule is implemented to allocate the capacity. The issue is the extent to which the supplier's allocation rule influences the supplier's profit, the retailer's profit and the supply chain's profit.

In this setting the firms that have the private information (the retailers) choose their actions simultaneously. Therefore, there is no information exchange among the firms. Even the supplier's capacity is fixed before the game starts, so the supplier is unable to use any information learned from the retailers' orders to choose a capacity. However, it is possible that correlation exists in the retailers' demand information, i.e., if a retailer observes his demand type to be high, then he might assess the other retailers' demand type to be high as well (if there is a positive correlation). Roughly speaking, in a Bayesian game each player uses Bayes' rule to update his belief regarding the types of the other players. An equilibrium is then a set of strategies for each type that is optimal given the updated beliefs with that type and the actions of all other types. See Fudenberg and Tirole (1991) for more information on Bayesian games.

6. Summary and Opportunities

As has been noted in other reviews, Operations Management has been slow to adopt GT. But because SCM is an ideal candidate for GT applications, we have recently witnessed an explosion of GT papers in SCM. As our survey indicates, most of these papers utilize only a few GT concepts, in particular the concepts related to non-cooperative static games. Some attention has been given to stochastic games but several other im-

portant areas need additional work: cooperative, repeated, differential, signaling, screening and Bayesian games.

The relative lack of GT applications in SCM can be partially attributed to the absence of GT courses from the curriculum of most doctoral programs in operations research/management. One of our hopes with this survey is to spur some interest in GT tools by demonstrating that they are intuitive and easy to apply for a person with traditional operations research training.

With the invention of the Internet, certain GT tools have received significant attention: web auctions gave a boost to auction theory, and numerous web sites offer an opportunity to haggle, thus making bargaining theory fashionable. In addition, the advent of relatively cheap information technology has reduced transaction costs and enabled a level of disintermediation that could not be achieved before. Hence, it can only become more important to understand the interactions among independent agents within and across firms. While the application of game theory to supply chain management is still in its infancy, much more progress will soon come.

Notes

1. This is an invited chapter for the book "Supply Chain Analysis in the eBusiness Era" edited by David Simchi-Levi, S. David Wu and Zuo-Jun (Max) Shen, to be published by Kluwer. <http://www.ise.ufl.edu/shen/handbook/>. Helpful comments from Gus Stuart and an anonymous referee are gratefully acknowledged.

2. Some may argue that GT should be a tool for choosing how a manager should play a game, which may involve playing against rational or semi-rational players. In some sense there is no conflict between these descriptive and normative roles for GT, but this philosophical issue surely requires more in-depth treatment than can be afforded here.

3. However, an argument can also be made that to predict rational behavior by players it is sufficient that players not choose dominated strategies, where a dominated strategy is one that yields a lower payoff than some other strategy (or convex combination of other strategies) for all possible strategy choices by the other players.

4. Strategy space is compact if it is closed and bounded.

5. Interior equilibrium is the one in which first-order conditions hold for each player. The alternative is boundary equilibrium in which at least one of the players select the strategy on the boundary of his strategy space.

6. The name is due to the fact that its source is unknown and dates back to 1960; Friedman (1986) was one of the first to treat Folk Theorem in detail.

7. A condition needed to insure attainability of an equilibrium solution is that the discount factor is large enough. The discount factor also affects effectiveness of trigger and many other strategies.

8. One interpretation of the value function is that it is the minimum value a coalition can guarantee for itself assuming the other players take actions that are most damaging to the coalition. But that can be criticized as overly conservative.

9. Another potential problem is that the core might be very large. However, as Brandenburger and Stuart (2003) point out, this may happen for a good reason: to interpret such situations, one can think of competition as not having much force in the game, hence the division of value will largely depend on the intangibles involved.

References

- Anand, K., R. Anupindi and Y. Bassok. 2002. Strategic inventories in procurement contracts. Working Paper, University of Pennsylvania.
- Anupindi, R., Y. Bassok and E. Zemel. 2001. A general framework for the study of decentralized distribution systems. *Manufacturing & Service Operations Management*, Vol.3, 349-368.
- Aumann, R. J. 1959. Acceptable Points in General Cooperative N-Person Games, pp. 287-324 in “*Contributions to the Theory of Games*”, Volume IV, A. W. Tucker and R. D. Luce, editors. Princeton University Press.
- Baiman, S., S. Netessine and H. Kunreuther. 2003. Procurement in supply chains when the end-product exhibits the weakest link property. Working Paper, University of Pennsylvania.
- Basar, T. and G.J. Olsder. 1995. *Dynamic noncooperative game theory*. SIAM, Philadelphia.
- Bernstein, F. and A. Federgruen. 2000. Comparative statics, strategic complements and substitute in oligopolies. Forthcoming in *Journal of Mathematical Economics*.
- Bernstein, F. and A. Federgruen. 2001a. Decentralized supply chains with competing retailers under Demand Uncertainty. Forthcoming in *Management Science*.
- Bernstein, F. and A. Federgruen. 2003. Pricing and Replenishment Strategies in a Distribution System with Competing Retailers. *Operations Research*, Vol.51, No.3, 409-426.
- Bernstein, F. and A. Federgruen. 2001c. A General Equilibrium Model for Decentralized Supply Chains with Price- and Service- Competition. Working Paper, Duke University.
- Bernstein, F. and A. Federgruen. 2002. Dynamic inventory and pricing models for competing retailers. Forthcoming in *Naval Research Logistics*.
- Bertsekas, D. P. 1999. *Nonlinear Programming*. Athena Scientific.
- Border, K.C. 1999. *Fixed point theorems with applications to economics and game theory*. Cambridge University Press.

- Brandenburger, A. and H.W. Stuart, Jr. 1996. Value-based business strategy. *Journal of Economics and Management Strategy*, Vol.5, No.1, 5-24.
- Brandenburger, A. and H.W. Stuart, Jr. 2003. Biform games. Working Paper, Columbia University.
- Cachon, G. 2001. Stock wars: inventory competition in a two-echelon supply chain. *Operations Research*, Vol.49, 658-674.
- Cachon, G. 2002. Supply chain coordination with contracts. Forthcoming in “*Handbooks in Operations Research and Management Science: Supply Chain management*”, S. Graves and T. de Kok, editors.
- Cachon, G. 2003. The allocation of inventory risk in a supply chain: push, pull and advanced purchase discount contracts. University of Pennsylvania working paper. Conditionally accepted, *Management Science*.
- Cachon, G. and C. Camerer. 1996. Loss avoidance and forward induction in coordination games. *Quarterly Journal of Economics*. Vol.112, 165-194.
- Cachon, G. and P. T. Harker. 2002. Competition and outsourcing with scale economies. *Management Science*, Vol.48, 1314-1333.
- Cachon, G. and M. Lariviere. 1999. Capacity choice and allocation: strategic behavior and supply chain performance. *Management Science*. Vol.45, 1091-1108.
- Cachon, G. and M. Lariviere. 2001. Contracting to assure supply: how to share demand forecasts in a supply chain. *Management Science*. Vol.47, 629-646.
- Cachon, G.P. and P.H. Zipkin. 1999. Competitive and cooperative inventory policies in a two-stage supply chain. *Management Science*, Vol.45, 936-953.
- Cachon, G. and G. Kok. 2002. Heuristic equilibrium in the newsvendor model with clearance pricing. Working Paper, University of Pennsylvania.
- Chen, F. 2001. Auctioning supply contracts. Working paper, Columbia University.
- Ren, J., Cohen, M., T. Ho., and C. Terwiesch. 2003. Sharing forecast information in a long-term supply chain relationship. Working Paper, University of Pennsylvania.
- Corbett, C. 2001. Stochastic inventory systems in a supply chain with asymmetric information: cycle stocks, safety stocks, and consignment stock. *Operations Research*. Vol.49, 487-500.
- Corbett C. J. and G. A. DeCroix. 2001. Shared-Savings Contracts for Indirect Materials in Supply Chains: Channel Profits and Environmental Impacts. *Management Science*, Vol.47, 881-893.

- Debo, L. 1999. Repeatedly selling to an impatient newsvendor when demand fluctuates: a supergame framework for co-operation in a supply chain. Working Paper, Carnegie Mellon University.
- Debreu, D. 1952. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, Vol.38, 886-893.
- Desai, V.S. 1992. Marketing-production decisions under independent and integrated channel structures. *Annals of Operations Research*, Vol.34, 275-306.
- Desai, V.S. 1996. Interactions between members of a marketing-production channel under seasonal demand. *European Journal of Operational Research*, Vol.90, 115-141.
- Eliashberg, J. and A.P. Jeuland. 1986. The impact of competitive entry in a developing market upon dynamic pricing strategies. *Marketing Science*, Vol.5, 20-36.
- Eliashberg, J. and R. Steinberg. 1987. Marketing-production decisions in an industrial channel of distribution. *Management Science*, Vol.33, 981-1000.
- Erhun F., P. Keskinocak and S. Tayur. 2000. Analysis of capacity reservation and spot purchase under horizontal competition. Working Paper, Georgia Institute of Technology.
- Feichtinger, G. and S. Jorgensen. 1983. Differential game models in management science. *European Journal of Operational Research*, Vol.14, 137-155.
- Filar, J. and K. Vrieze. 1996. *Competitive Markov decision processes*. Springer-Verlag.
- Friedman, J.W. 1986. *Game theory with applications to economics*. Oxford University Press.
- Fudenberg, D. and J. Tirole. 1991. *Game theory*. MIT Press.
- Gaimon, C. 1989. Dynamic game results of the acquisition of new technology. *Operations Research*, Vol.3, 410-425.
- Gale, D. and H. Nikaido. 1965. The Jacobian matrix and global univalence of mappings. *Mathematische Annalen*, Vol.159, 81-93.
- Gillemin V. and A. Pollak. 1974. *Differential Topology*. Prentice Hall, NJ.
- Granot, D. and G. Sosic. 2001. A three-stage model for a decentralized distribution system of retailers. Forthcoming, *Operations Research*.
- Hall, J. and E. Porteus. 2000. Customer service competition in capacitated systems. *Manufacturing & Service Operations Management*, Vol.2, 144-165.
- Hartman, B. C., M. Dror and M. Shaked. 2000. Cores of inventory centralization games. *Games and Economic Behavior*, Vol.31, 26-49.

- Heyman, D. P. and M. J. Sobel. 1984. *Stochastic models in Operations Research*, Vol.II: Stochastic Optimization. McGraw-Hill.
- Horn, R.A. and C.R. Johnson. 1996. *Matrix analysis*. Cambridge University Press.
- Kamien, M.I. and N.L. Schwartz. 2000. *Dynamic optimization: the calculus of variations and optimal control in economics and management*. North-Holland.
- Kirman, A.P. and M.J. Sobel. 1974. Dynamic oligopoly with inventories. *Econometrica*, Vol.42, 279-287.
- Kreps, D. M. 1990. *A Course in Microeconomic Theory*. Princeton University Press.
- Kreps, D. and R. Wilson. 1982. Sequential equilibria. *Econometrica*, Vol.50, 863-894.
- Kuhn, H. W. 1953. Extensive Games and the Problem of Information. pp. 193-216 in “*Contributions to the Theory of Games*”, Volume II, H. W. Kuhn and A. W. Tucker, editors. Princeton University Press.
- Lal, R. 1990. Price promotions: limiting competitive encroachment. *Marketing Science*, Vol.9, 247-262.
- Lariviere, M.A. and E.L. Porteus. 2001. Selling to the newsvendor: an analysis of price-only contracts. *Manufacturing & Service Operations Management*, Vol.3, 293-305.
- Lederer, P. and L. Li. 1997. Pricing, production, scheduling, and delivery-time competition. *Operations Research*, Vol.45, 407-420.
- Li, L. and S. Whang. 2001. Game theory models in operations management and information systems. In “*Game theory and business applications*”, K. Chatterjee and W.F. Samuelson, editors. Kluwer Academic Publishers.
- Lippman, S.A. and K.F. McCardle. 1997. The competitive newsboy. *Operations Research*, Vol.45, 54-65.
- Lucas, W.F. 1971. An overview of the mathematical theory of games. *Management Science*, Vol.18, 3-19.
- Mahajan, S and G. van Ryzin. 1999a. Inventory competition under dynamic consumer choice. *Operations Research*, Vol.49, 646-657.
- Mahajan, S and G. van Ryzin. 1999b. Supply chain coordination under horizontal competition. Working Paper, Columbia University.
- Majumder, P. and H. Groenevelt. 2001a. Competition in remanufacturing. *Production and Operations Management*, Vol.10, 125-141.
- Majumder, P. and H. Groenevelt. 2001b. Procurement competition in remanufacturing. Working Paper, Duke University.
- Moulin, H. 1986. *Game theory for the social sciences*. New York University Press.

- Moulin, H. 1995. *Cooperative microeconomics: a game-theoretic introduction*. Princeton University Press.
- Muller, A., M. Scarsini and M. Shaked. 2002. The newsvendor game has a nonempty core. *Games and Economic Behavior*, Vol.38, 118-126.
- Mukhopadhyay, S.K. and P. Kouvelis. 1997. A differential game theoretic model for duopolistic competition on design quality. *Operations Research*, Vol.45, 886-893.
- Myerson, R.B. 1997. *Game theory*. Harvard University Press.
- Nash, J. F. 1950. Equilibrium Points in N-Person Games. *Proceedings of the National Academy of Sciences of the United States of America*, Vol.36, 48-49.
- Netessine, S. and N. Rudi. 2001a. Supply chain structures on the Internet and the role of marketing-operations interaction. Forthcoming in "Supply chain analysis in e-business era", D. Simchi-Levi, S.D. Wu and M. Shen, Edts., Kluwer.
- Netessine, S. and N. Rudi. 2001b. Supply Chain choice on the Internet. Working Paper, University of Pennsylvania. Available at <http://www.netessine.com>.
- Netessine, S. and N. Rudi. 2003. Centralized and competitive inventory models with demand substitution. *Operations Research*, Vol.53, 329-335.
- Netessine, S. and R. Shumsky. 2001. Revenue management games: horizontal and vertical competition. Working Paper, University of Pennsylvania. Available at <http://www.netessine.com>.
- Netessine, S., N. Rudi and Y. Wang. 2002. Dynamic inventory competition and customer retention. Working Paper, University of Pennsylvania, available at <http://www.netessine.com>.
- Netessine, S. and F. Zhang. 2003. The impact of supply-side externalities among downstream firms on supply chain efficiency. Working Paper, University of Pennsylvania, available at <http://www.netessine.com>.
- Parlar, M. 1988. Game theoretic analysis of the substitutable product inventory problem with random demands. *Naval Research Logistics*, Vol.35, 397-409.
- Plambeck, E. and T. Taylor. 2001a. Sell the plant? The impact of contract manufacturing on innovation, capacity and profitability. Working Paper, Stanford University.
- Plambeck, E. and T. Taylor. 2001b. Renegotiation of supply contracts. Working Paper, Stanford University.
- Porteus, E. and S. Whang. 1999. Supply chain contracting: non-recurring engineering charge, minimum order quantity, and boilerplate contracts. Working paper, Stanford University.

- Rosen, J.B. 1965. Existence and uniqueness of equilibrium points for concave N-person games. *Econometrica*, Vol.33, 520-533.
- Rudi, N., S. Kapur and D. Pyke. 2001. A two-location inventory model with transshipment and local decision making. *Management Science*, Vol.47, 1668-1680.
- Selten, R. 1965. Spieltheoretische behandlung eine oligopolmodells mit nachfragetragheit. *Zeitschrift fur die gesamte staatswissenschaft*. Vol.12, 301-324.
- Selten, R. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, Vol.4, 25-55.
- Shapley, L. 1953a. Stochastic games. *Proceedings of the National Academy of Sciences*, Vol.39, 1095-1100.
- Shapley, L. 1953b. A value for n -person game. pp. 307-317 in “*Contributions to the Theory of Games*”, Volume II, H. W. Kuhn and A. W. Tucker, editors. Princeton University Press.
- Shubik, M. 1962. Incentives, decentralized control, the assignment of joint costs and internal pricing. *Management Science*, Vol.8, 325-343.
- Shubik, M. 2002. Game theory and operations research: some musings 50 years later. *Operations Research*, Vol.50, 192-196.
- Sobel, M.J. 1971. Noncooperative stochastic games. *Annals of Mathematical Statistics*, Vol.42, 1930-1935.
- Stackelberg, H. von. 1934. *Markform and Gleichgewicht*. Vienna: Julius Springer.
- Stidham, S. 1992. Pricing and capacity decisions for a service facility: stability and multiple local optima. *Management Science*. Vol.38, 1121-1139.
- Stuart, H. W., Jr. 2001. Cooperative games and business strategy. In “*Game theory and business applications*”, K. Chatterjee and W.F. Samuelson, editors. Kluwer Academic Publishers.
- Tarski, A. 1955. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, Vol.5, 285-308.
- Taylor, T. A. and E. L. Plambeck. 2003. Supply chain relationships and contracts: the impact of repeated interaction on capacity investment and procurement. Working paper, Columbia University.
- Tayur, S. and W. Yang. 2002. Equilibrium analysis of a natural gas supply chain. Working Paper, Carnegie Mellon University.
- Topkis, D. M. 1998. *Supermodularity and complementarity*. Princeton University Press.
- van Mieghem, J. 1999. Coordinating investment, production and sub-contracting. *Management Science*, Vol.45, 954-971.

- van Mieghem, J. and M. Dada. 1999. Price versus production postponement: capacity and competition. *Management Science*, Vol.45, 1631-1649.
- Varian, H. 1980. A model of sales. *American Economic Review*, Vol.70, 651-659.
- Vickrey W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, Vol.16, 8-37.
- Vives, X. 1999. *Oligopoly pricing: old ideas and new tools*. MIT Press.
- von Neumann, J. and O. Morgenstern. 1944. *Theory of games and economic behavior*. Princeton University Press.
- Wang, Y. and Y. Gerchak. 2003. Capacity games in assembly systems with uncertain demand. *Manufacturing & Service Operations Management*, Vol.5, No.3, 252-267.
- Wang, Q. and M. Parlar. 1989. Static game theory models and their applications in management science. *European Journal of Operational Research*, Vol.42, 1-21.
- Wang, Q. and M. Parlar. 1994. A three-person game theory model arising in stochastic inventory control theory. *European Journal of Operational Research*, Vol.76, 83-97.