1

Introduction

1.1 Historical perspective

Dynamic game theory brings together four features that are key to many situations in economics, ecology and elsewhere: optimizing behavior, the presence of multiple agents/players, enduring consequences of decisions and robustness with respect to variability in the environment.

To deal with problems which have these four features the dynamic game theory methodology splits the modeling of the problem into three parts. One part is the modeling of the environment in which the agents act. To obtain a mathematical model of the agents' environment a set of differential or difference equations is usually specified. These equations are assumed to capture the main dynamical features of the environment. A characteristic property of this specification is that these dynamic equations mostly contain a set of so-called 'input' functions. These input functions model the effect of the actions taken by the agents on the environment during the course of the game. In particular, by viewing 'nature' as a separate player in the game who can choose an input function that works against the other player(s), one can model worst-case scenarios and, consequently, analyze the robustness of the 'undisturbed' game solution.

A second part is the modeling of the agents' objectives. Usually the agents' objectives are formalized as cost/utility functions which have to be minimized. Since this minimization has to be performed subject to the specified dynamic model of the environment, techniques developed in optimal control theory play an important role in solving dynamic games. In fact, from a historical perspective, the theory of dynamic games arose from the merging of static game theory and optimal control theory. However, this merging cannot be done without further reflection. This is exactly what the third modeling part is about. To understand this point it is good to summarize the rudiments of static games.

Most research in the field of static game theory has been – and is being – concentrated on the normal form of a game. In this form all possible sequences of decisions of each player are set out against each other. So, for example, for a two-player game this results in a matrix structure. Characteristic for such a game is that it takes place in one moment of

COPYRIGHTED MATERIAL

LQ Dynamic Optimization and Differential Games J. Engwerda
© 2005 John Wiley & Sons, Ltd

time: all players make their choice once and simultaneously and, dependent on the choices made, each player receives his payoff. In such a formulation important issues like the order of play in the decision process, information available to the players at the time of their decisions, and the evolution of the game are suppressed, and this is the reason this branch of game theory is usually classified as 'static'. In case the agents act in a dynamic environment these issues are, however, crucial and need to be properly specified before one can infer what the outcome of the game will be. This specification is the third modeling part that characterizes the dynamic game theory methodology.

In this book we study a special class of dynamic games. We study games where the environment can be modeled by a set of linear differential equations and the objectives can be modeled as functions containing just affine quadratic terms. Concerning the information structure of the game we will basically describe two cases: the 'open-loop' and the 'linear feedback' case. A proper introduction of these notions is postponed until the relevant chapters later on.

The popularity of these so-called linear quadratic differential games is caused on the one hand by practical considerations. To some extent these kinds of differential games are analytically and numerically solvable. If one leaves this track, one easily gets involved in the problem of solving sets of nonlinear partial differential equations, and not many of these equations can be solved analytically. Even worse, when the number of state variables is more than two in these equations, a numerical solution is in general hard to obtain. On the other hand this linear quadratic problem setting naturally appears if the agents' objective is to minimize the effect of a small perturbation of their nonlinear optimally controlled environment. By solving a linear quadratic control problem, and using the optimal actions implied by this problem, players can avoid most of the additional cost incurred by this perturbation (section 5.1).

So, linear quadratic differential games are a subclass of dynamic games. As already indicated above, optimal control techniques play an important role in solving dynamic games and, in fact, optimal control theory is one of the roots of dynamic game theory. For these reasons the first part of this book (Chapters 2–5) gives, broadly speaking, an introduction to the basics of the theory of dynamic optimization. To appreciate this theory we next provide a short historical overview of its development.

To outline the field, optimal control theory is defined as the subject of obtaining optimal (i.e. minimizing or maximizing) solutions and developing numerical algorithms for one-person single-objective dynamic decision problems.

Probably the first recorded feedback control application is the water clock invented by the Greek Ktesibios around 300 BC in Alexandria, Egypt. This was definitely a successful design as similar clocks were still used around 1260 AD in Baghdad. The theory on optimal control has its roots in the calculus of variations. The Greek Pappus of Alexandria[1] already posed 300 AD the isoperimetric problem, i.e. to find a closed plane curve of a given length which encloses the largest area, and concluded that this was a circle. Remarkably, the most essential contribution towards its rigorous proof was only given in 1841 by Steiner[2]. Some noteworthy landmarks in between are the derivation by

[1]Pappus, ± 290–± 350, born in Alexandria (Egypt), was the last of the great Greek geometers. He is sometimes called the founding father of projective geometry.

[2]Steiner, 1796–1863, was a Swiss mathematician who first went to school at the age of 18. He made very significant contributions to projective geometry.

Fermat[3] around 1655 of the sine law of refraction (as proposed by Snell[4]) using the principle that light always follows the shortest possible path; the publication of Newton's Principia[5] in 1687 in which he analyses the motion of bodies in resisting and non-resisting media under the action of centripetal forces; and the formulation and solution in 1696 of the 'Brachistochrone' problem by Johann Bernoulli[6], i.e. to find the curve along which a particle uses the minimal time to slide between two points A and B, if the particle is influenced by gravitational forces only – this curve turns out to be a cycloid. The name calculus of variation was introduced by Euler[7] in 1766 and both he (in 1744) and Lagrange[8] (1759, 1762 and 1766) contributed substantially to the development of this theory. Both their names are attached to the famous Euler–Lagrange differential equation that an optimal function and its derivative have to satisfy in order to solve a dynamic optimization problem. This theory was further developed in the nineteenth century by Hamilton[9] in his papers on general methods in dynamics (1834,1835), and in the lectures by Weierstrass[10] at the University of Berlin during 1860–1890 and Jacobi[11] who carried out important research in partial differential equations and used this theory to analyze equations describing the motion of dynamical systems. Probably the first mathematical model to describe plant behavior for control purposes is due to Maxwell[12], who in 1868 used differential equations to explain instability problems encountered with James Watt's flyball governor (later on used to regulate the speed of steam-engine vehicles). 'Forerunners' of modern optimal control theory, associated with the maximum principle, are Valentine (1937), McShane[13] (1939), Ambartsumian (1943) and Hesteness (1949) (the expanded version of this work was later published in 1966). Particularly in the 1930s and

[3]Fermat, 1601–1665, was a French lawyer who is famous for his mathematical contributions to the algebraic approach to geometry and number theory.

[4]Snell, 1580–1626, was a Dutch mathematician/lawyer who made significant contributions to geodesy and geometric optics.

[5]Newton, 1643–1727, born in England, was famous for the contributions he made in the first half of his career to mathematics, optics, physics and astronomy. The second half of his life he spent as a government official in London.

[6]Johann Bernoulli, 1667–1748, was a Swiss mathematician who made significant contributions to analysis.

[7]Euler, 1707–1783, a Swiss mathematician/physician was one of the most prolific writers on mathematics of all time. He made decisive and formalistic contributions to geometry, calculus, number theory and analytical mechanics.

[8]Lagrange, 1736–1813, was a Sardinian mathematician who made substantial contributions in various areas of physics, the foundation of calculus, dynamics, probability and number theory.

[9]Hamilton, 1805–1865, an Irish mathematician/physician, introduced the characteristic function and used this to study dynamics. Moreover, he introduced and studied the algebra of quaternions that play an important role in mathematical physics.

[10]Weierstrass, 1815–1897, was a German mathematician. The standards of rigour Weierstrass set in his courses strongly affected the mathematical world and for that reason he is sometimes called the father of modern analysis.

[11]Jacobi, 1804–1851, was a German mathematician, who was prepared (but not allowed) to enter university when he was 12. He was a very prolific writer in many fields of mathematics. His contributions to the theory of partial differential equations and determinants are well-known.

[12]Maxwell, 1831–1879, a Scottish mathematician/physicist, published his first paper when he was only 14. His most well-known contributions are about electricity and magnetism and his kinetic theory of gases.

[13]McShane, 1904–1989, was an American mathematician, well-known for his work in the calculus of variations, ballistics, integration theory and stochastic differential equations.

1940s frequency domain methods and Laplace transformation techniques were used to study control problems. At the beginning of the second half of the twentieth century further theoretical impetus was provided on the one hand by the Russian Pontrjagin[14], around 1956, with the development of the so-called maximum principle. This culminated in the publication of his book (1961). On the other hand, around 1950, the Americans Bellman and Isaacs started the development of the dynamic programming principle[15] which led to the publication of the books by Bellman (1956) and Bellman and Dreyfus (1962).

Furthermore, it was recognized in the late 1950s that a state space approach could be a powerful tool for the solution of, in particular, linear feedback control problems. The main characteristics of this approach are the modeling of systems using a state space description, optimization in terms of quadratic performance criteria, and incorporation of Kalman–Bucy optimal state reconstruction theory. The significant advantage of this approach is its applicability to control problems involving multi-input multi-output systems and time-varying situations (for example Kwakernaak and Sivan, 1972). A historical overview on early control theory can be found in Neustadt (1976).

Progress in stochastic, robust and adaptive control methods from the 1960s onwards, together with the development of computer technology, have made it possible to control much more accurately dynamical systems which are significantly more complex. In Bushnell (1996), one can find a number of papers, including many references, concerning the more recent history of control theory.

For more or less the same reasons we covered the historical background of optimal control theory we next trace back some highlights of game theory, ending the overview more or less at the time the theory of dynamic games emerged. This avoids the challenge of providing an overview of its most important recent theoretical developments.

The first static cooperative game problem reported seems to date back to 0–500 AD. In the Babylonian Talmud, which serves as the basis of Jewish religious, criminal and civil law, the so-called marriage contract problem is discussed. In this problem it is specified that when a man who has three wives dies, wives receive 100, 200 and 300, respectively. However, it also states that if the estate is only worth 100 all three receive the same amount, if it is worth 200 they receive 50, 75 and 75, respectively, and if it is worth 300 they receive a proportional amount, i.e. 50, 100 and 150, respectively. This problem puzzled Talmudic scholars for two millennia. It was only recognized in 1985 that the solution presented by the Talmud can be interpreted using the theory of co-operative games. Some landmarks in this theory are a book on probability theory written by

[14]Pontrjagin, 1908–1988, a Russian mathematician who due to an accident was left blind at 14. His mother devoted herself to help him succeed to become a mathematician. She worked for years in fact as his private secretary, reading scientific works aloud to him, writing in the formulae in his manuscripts, correcting his work, though she had no mathematical training and had to learn to read foreign languages. He made important contributions to topology, algebra and, later on, control theory.

[15]Both scientists worked in the late 1940s and early 1950s at the Research and New Development (RAND) Corporation in Santa Monica, California, USA. During presentations and discussions at various seminars held at RAND at that time the dynamic programming principle probably arose as a principle to solve dynamic optimization problems. From the discussions later on (Breitner, 2002) it never became clear whether just one or both scientists should be considered as the founding father(s) of this principle.

Montmort[16] in 1708 in which he deals in a systematic way with games of chance. In 1713 Waldegrave, inspired by a card game, provided the first known minimax mixed strategy solution to a two-person game. Cournot[17] (1838) discusses for the first time the question of what equilibrium price might result in the case of two producers who sell an identical product (duopoly). He utilizes a solution concept that is a restricted version of the non-cooperative Nash equilibrium. The theory of cooperation between economic agents takes its origin from economic analysis. Probably Edgeworth[18] and Pareto[19] provided the first definitions of a cooperative outcome. Edgeworth (1881) proposed the contract curve as a solution to the problem of determining the outcome of trading between individuals, whereas Pareto (1896) introduced the notion of efficient allocation. Both used the formalism of ordinal utility theory. Zermelo[20] (1913) presented the first theorem of game theory which asserts that chess is a strictly determined game, i.e. assuming that only an a priori fixed number of moves is allowed, either (i) white has a strategy which always wins; or (ii) white has a strategy which always at least draws, but no strategy as in (i); or (iii) black has a strategy which always wins. Fortunately, no one knows which of the above is actually the case. There are games where the assumption that each player may choose from only a finite number of actions seems to be inappropriate. Consider, for example, the 'princess and the monster' game (see Foreman (1977) and Başar and Olsder (1999)). In this game, which is played in a completely dark room, there is a monster who wants to catch a princess. The moment both bump into each other the game terminates. The monster likes to catch the princess as soon as possible, whereas the princess likes to avoid the monster as long as possible. In this game the optimal strategies cannot be deterministic. For, if the monster were to have a deterministic optimal strategy, then the princess would be able to calculate this strategy. This would enable her to determine the monster's path and thereby to choose for herself a strategy such that she avoids the monster forever. Therefore, an optimal strategy for the monster (if it exists) should have random actions, so that his strategy cannot be predicted by the princess. Such a strategy is called mixed. Borel[21] published from 1921–1927 five notes in which he gave the first modern formulation of a mixed strategy along with finding the minimax[22] solution for

[16]Montmort, 1678–1719, was a French mathematician.

[17]Cournot, 1801–1877, a French mathematician was the pioneer of mathematical economics.

[18]Edgeworth, 1845-1926, an Irish economist/mathematician was well-known for his work on utility theory and statistics.

[19]Pareto, 1848–1923, was an Italian economist/sociologist who studied classics and engineering in Turin. After his studies in 1870 he worked for a couple of years in industry. From 1889 onwards he began writing numerous polemical articles against the Italian government and started giving public lectures. In 1893 he succeeded Walras at the University of Lausanne (Switzerland). Famous are his Manual of Political Economy (1906), introducing modern microeconomics, and his Trattato di Sociologia Generale (1916), explaining how human action can neatly be reduced to residue (non-logical sentiments) and derivation (afterwards justifications).

[20]Zermelo, 1871–1953, was a German mathematician/physician famous for his work on axiomatic set theory.

[21]Borel, 1871–1956, was a French mathematician/politician famous for his work on measure theory. In the second half of his life he embarked on a political career and became Minister of the Navy.

[22]The idea of a minimax solution is that a player wants to choose that action which minimizes the maximum risk that can occur due to the actions of his opponent(s).

two-person games with three or five possible strategies. Von Neumann[23] (1928) considered two-person zero-sum games. These are games where the revenues of one player are costs for the other player. He proved that every two-person zero-sum game with finitely many pure strategies for each player is determined and introduces the extensive form of a game. The extensive form of a game basically involves a tree structure with several nodes and branches, providing an explicit description of the order of play and the information available to each player at the time of his decision. The game evolves from the top of the tree to the tip of one of its branches. The case that players decide to cooperate in order to maximize their profits was considered by Zeuthen[24] (1930). If players agree to cooperate in order to maximize their profits the question arises as to what extent efforts will be used to maximize the individual profit of each single player. This is called the bargaining problem. Zeuthen (1930) proposes a solution to the bargaining problem which Harsanyi[25] (1956) showed is equivalent to Nash's bargaining solution. A distinction between the order of play was introduced by von Stackelberg[26] (1934) within the context of economic competition. He distinguished between the first mover, called the leader, and the second mover, called the follower. The idea is that the follower can observe the move of the leader and subsequently act. The seminal work of Von Neumann and Morgenstern[27] (1944) presents the two-person zero-sum theory; the notion of a cooperative game with transferable utility, its coalitional form and stable sets; and axiomatic utility theory. In their book they argue that economics problems may be analyzed as games. Once all irrelevant details are stripped away from an economics problem, one is left with an abstract decision problem – a game. The book led to an era of intensive game theory research. The next cornerstone in static game theory was set by Nash from 1950–1953 in four papers. He proved the existence of a strategic equilibrium for non-cooperative games – the Nash equilibrium – and proposed the 'Nash program', in which he suggested approaching the study of cooperative games via their reduction to non-cooperative form (1950a,1951). A Nash equilibrium was defined as a strategy combination – consisting of one stategy for each player – with the property that no player can gain (in terms of utility) by unilaterally deviating from it. Hence this equilibrium solution is self-enforcing. That is, it is an optimal solution for each player as long as his opponent players stick to their recommendations. Unfortunately, it turns out that a game may have more than one such Nash equilibrium. Since not all equilibria are in general equally attractive, refinement criteria for selecting among multiple equilibria were proposed later on (for example van Damme, 1991). Another, still relevant, issue associated with this non-uniqueness is how one can determine numerically all refined Nash equilibria (for example Peeters, 2002). In his two papers on bargaining theory (Nash, 1950b,1953), Nash founded axiomatic bargaining theory, proved the existence of the Nash bargaining solution and provided the first execution of the Nash program. In

[23]Von Neumann, 1903–1957, was a Hungarian mathematician who was a pioneer in various fields including quantum mechanics, algebra, applied mathematics and computer science.

[24]Zeuthen, 1888–1959, was a Danish economist known for his work on general equilibrium theory, bargaining and monopolistic competition.

[25]Harsanyi, 1920–2000, was a Hungarian pharmacist/philosopher/economist well-known for his contributions to game theory.

[26]Von Stackelberg, 1905–1946, was a German economist.

[27]Morgenstern, 1902–1976, was an Austrian economist.

axiomatic bargaining theory one tries to identify rules that yield a fair sharing of the benefits of cooperation (see also Chapter 6). From this time onwards the theoretical developments in game theory rapidly increased. By consulting the latest textbooks in this field (for example Tijs (1981) or Fudenberg and Tirole (1991)) or the website of Walker, (2001) one can get an impression of these developments.

Probably as a spin-off of all these new ideas in control and game theory in the late 1940s and early 1950s the first dynamic game models came to light. The initial drive towards the development of this theory was provided by Isaacs during his stay at the RAND Corporation from 1948–1955. He wrote the first paper on games of pursuit (Isaacs, 1951). In this paper the main ideas for the solution of two-player, zero-sum dynamic games of the pursuit–evasion type are already present. He furthered these ideas (1954–1955) and laid a basis for the theory of dynamic games within the framework of two-person zero-sum games. This theory was first discussed by Berkovitz (1961). The first contribution on nonzero-sum games seems to date back to Case (1967). A nice historic overview on the early years of differential games and Isaacs' contributions to this field can be found in Breitner (2002).

The first official papers on dynamic games were published in a special issue of the Annals of Mathematics Studies edited by Dresher, Tucker and Wolfe (1957). In a special section of this volume entitled 'Games with a continuum of moves' Berkovitz, Fleming and Scarf published their papers (Berkovitz and Fleming, 1957; Fleming, 1957; Scarf, 1957). On the Russian side early official contributions on dynamic games have been published by Kelendzeridze (1961), Petrosjan (1965), Pontrjagin (1961) and Zelikin and Tynyanskij (1965). The books written by Isaacs (1965) and Blaquiere, Gerard and Leitmann (1969) document the theoretical developments of dynamic game theory during its first two decades. The historical development of the theory since the late 1960s is documented by the works of, for example, Friedman (1971), Leitmann (1974), Krasovskii and Subbotin (1988), Mehlmann (1988), Başar and Olsder (1999) and Haurie (2001). Particularly in Başar and Olsder (1999) one can find at the end of each chapter a section where relevant historical remarks are included concerning the subjects discussed in that chapter.

Current applications of differential games range from economics, financial engineering, ecology and marketing to the military. The work of Dockner *et al.* (2000) provides an excellent comprehensive, self-contained survey of the theory and applications of differential games in economics and management science. The proceedings and the associated Annals of the International Symposia on Dynamic Games and Applications held every other year (for example Petrosjan and Zenkevich, 2002) document the development of both theory and applications over the last 20 years.

We conclude this section by presenting a historical outline of the development of the theory on non-cooperative linear quadratic differential games. As already indicated the linear quadratic differential games constitute a subclass of differential games. Starr and Ho might be called the founding fathers of this theory. Ho, Bryson and Baron (1965) analyzed the particular class of pursuit–evasion games, and the results were later put into a rigorous framework by Schmitendorf (1970). With their paper, Starr and Ho (1969) generalized the zero-sum theory developed by Isaacs. Using the Hamilton–Jacobi theory they provided a sufficient condition for existence of a linear feedback Nash equilibrium for a finite-planning horizon.

Lukes (1971) showed that, if the planning horizon in the game is chosen to be sufficiently small, the game always has – for every initial state of the system a unique

linear feedback Nash equilibrium. Moreover, this equilibrium can be computed by solving a set of so-called feedback Nash Riccati differential equations. Papavassilopoulos and Cruz (1979) show that if the set of strategy spaces is restricted to analytic functions of the current state and time, then the Nash equilibrium is unique, if it exists. Bernhard (1979) considers the zero-sum game with the additional restriction that the final state should lie in some prespecified linear subspace. Mageirou (1976) considered the infinite-horizon zero-sum game. An important point demonstrated by this paper and made more explicit by Jacobson (1977) is that the strategy spaces should be clearly defined before one can derive equilibria. Papavassilopoulos, Medanić and Cruz (1979) discussed parametric conditions under which the coupled set of algebraic feedback Nash Riccati equations has a solution. If these conditions are met the infinite-horizon game has at least one feedback Nash equilibrium. Papavassilopoulos and Olsder (1984) demonstrate that an infinite-horizon game may have either none, a unique or multiple feedback Nash equilibria even though every finite-horizon version of it has a unique feedback Nash equilibrium. In particular they present a sufficient condition under which the set of feedback Nash Riccati differential equations has a solution. This last result was generalized by Freiling, Jank and Abou-Kandil (1996). Weeren, Schumacher and Engwerda (1999) give an asymptotic analysis of the regular finite-planning two-player scalar game. They show that this game always has a unique equilibrium but that the convergence of the equilibrium actions depends on the scrap value. Three different convergence schemes may occur and the equilibrium actions always converge in this regular case to a solution of the infinite-horizon game. The number of equilibria for the scalar N-player infinite-horizon game was studied in Engwerda (2000b). For the two-player case parameter conditions under which this game has a unique equilibrium were derived in Lockwood (1996) and Engwerda (2000a).

The problem of calculating the solutions of the feedback Riccati differential equations is adressed in Cruz and Chen (1971) and Ozgüner and Perkins (1977). Iterative algorithms to calculate a stabilizing solution of the algebraic feedback Nash equations were developed by Krikelis and Rekasius (1971), Tabak (1975), Mageirou (1977) and Li and Gajic (1994). A disadvantage of these algorithms is that they depend on finding good initial conditions and provide just one solution (if they converge) of the equations. In Engwerda (2003) an algorithm based on determining the eigenstructure of a certain matrix was presented to calculate the whole set of stabilizing solutions in case the system is scalar.

Under the assumption that the planning horizon is not too long, Friedman (1971) showed that the game will have a unique open-loop Nash equilibrium (see also Starr (1969)). For an arbitrary finite-planning horizon length Lukes and Russel (1971) presented a sufficient condition for the existence and uniqueness of open-loop equilibria in terms of the invertibility of a Hilbert space operator. Eisele (1982) used this latter approach to show that if this operator is not invertible the game either has no open-loop equilibrium solution or an infinite number of solutions, this depending on the initial state of the system. Feucht (1994) reconsidered the open-loop problem for a general indefinite cost function and studied in particular its relationship with the associated set of open-loop Riccati differential equations.

Analytic solutions of these Riccati differential equations have been studied in Simaan and Cruz (1973), Abou-Kandil and Bertrand (1986), Jódar and Abou-Kandil (1988,1989), Jódar (1990), Jódar and Navarro (1991a,b), Jódar, Navarro and Abou-Kandil (1991), and Abou-Kandil, Freiling and Jank (1993). These results were generalized by Feucht (1994).

The basic observation made in these references is that this set of differential equations constitute an ordinary (non-symmetric, high-order) Riccati differential equation and its solution (Reid, 1972) can thus be analyzed as the solution of a set of linear differential equations (see also Abou-Kandil *et al.*, 2003). If some additional parametric assumptions are made the set of coupled equations reduces to one single (non-symmetric, low-order) Riccati differential equation. An approximate solution is derived in Simaan and Cruz (1973) and Jódar and Abou-Kandil (1988). Scalzo (1974) showed that if the controls are constrained to take values in compact convex subsets the finite-horizon game always has an open-loop equilibrium. This is irrespective of the duration of the game. Engwerda and Weeren (1994) and Engwerda (1998a) studied both the limiting behavior of the finite-planning horizon solutions as the final time approaches infinity and the infinite-horizon case using a variational approach. An algorithm to calculate all equilibria for the infinite-horizon game was provided in Engwerda (1998b). Kremer (2002) used the Hilbert space approach to analyze the infinite-horizon game and showed in particular that similar conclusions hold in this case as those obtained by Eisele.

To model uncertainty basically two approaches have been taken in literature (but see also Bernhard and Bellec (1973) and Broek, Engwerda and Schumacher (2003) for a third approach). Usually either a stochastic or a worst-case approach is taken. A stochastic approach, for example, is taken in Kumar and Schuppen (1980), Başar (1981), Bagchi and Olsder (1981) and Başar and Li (1989). A worst-case approach (see, for example, the seminal work by Başar and Bernhard (1995)) is taken, in a cooperative setting, by Schmitendorf (1988). The non-cooperative open-loop setting is dealt with by Kun (2001) and the linear feedback setting by Broek, Engwerda and Schumacher (2003).

Applications in economics are reported in various fields. In, for example, in industrial organization by Fershtman and Kamien (1987), Reynolds (1987), Tsutsui and Mino (1990), Chintagunta (1993), Jørgensen and Zaccour (1999, 2003); in exhaustible and renewable resources by Hansen, Epple and Roberts (1985), Mäler and de Zeeuw (1998) and Zeeuw and van der Ploeg (1991); in interaction between monetary and fiscal authorities by Pindyck (1976), Kydland (1976), Hughes-Hallett (1984), Neese and Pindyck (1984), Tabellini (1986), Petit (1989), Hughes-Hallett and Petit (1990), van Aarle, Bovenberg and Raith (1995), Engwerda, van Aarle and Plasmans (1999, 2002) van Aarle *et al.* (2001) and van Aarle, Engwerda and Plasmans (2002); in international policy coordination by Miller and Salmon (1985a,b), Cohen and Michel (1988), Curie, Holtham and Hughes-Hallett (1989), Miller and Salmon (1990) and Neck and Dockner (1995); and in monetary policy games by Obstfeld (1991) and Lockwood and Philippopoulos (1994). In this context it should be mentioned that various discrete-time macroeconomic game models have been estimated and analyzed in literature (e.g. the SLIM model developed by Douven and Plasmans (1996) and the Optgame model developed by Neck *et al.* (2001).

In literature information structures different from the ones that are considered in this book have also been investigated. For instance Foley and Schmitendorf (1971) and Schmitendorf (1970) considered the case that one player has open-loop information and the other player uses a linear-feedback strategy. Furthermore Başar (1975,1977 or, for its discrete-time counterpart, 1974) showed that finite-planning horizon linear quadratic differential games also permit multiple nonlinear Nash equilibria if at least one of the players has access to the current and initial state of the system (the dynamic information case). Finally, we should point out the relationship between the optimal control of stochastic linear systems with an exponential performance criterion and zero-sum

differential games, which enables a stochastic interpretation of worst-case design of linear systems (Jacobson (1973) and Broek, Engwerda and Schumacher (2003a) or Klompstra (2000) in a discrete-time framework).

1.2 How to use this book

This book is a self-contained introduction to linear quadratic differential games. The book is introductory, but not elementary. It requires some basic knowledge of mathematical analysis, linear algebra and ordinary differential equations. The last chapter also assumes some elementary knowledge of probability theory. The topics covered in the various chapters can be followed up to a large extent in related literature. In particular the sections entitled 'Notes and references' which end each chapter can be regarded as pointers to the sources consulted and related items that are not mentioned in this book.

This book is intended to be used either as a textbook by students in the final year of their studies or as a reference work. The material is written in such a way that most of the material can be read without consulting the mathematical details. Lengthy proofs are provided, in most cases, in the appendix to each chapter. Broadly speaking, the book consists of two parts: the first part (Chapters 2–5) is about dynamic optimization with, as a special case, the linear quadratic control problem. The second part (Chapters 6–9) is about linear quadratic differential games. So, this book could be used to teach a first semester introductory course on dynamic optimization and a second semester course on linear quadratic differential games. Throughout this book the theory is illustrated by examples which are often taken from the field of economics. These examples should help the reader to understand the presented theory.

1.3 Outline of this book

A summary of each chapter is given below. Note that some of the statements in this section are not precise. They hold under certain assumptions which are not explicitly stated. Readers should consult the corresponding chapters for the exact results and conditions.

Chapter 2 reviews some basic linear algebra which in some instances goes beyond the introductory level. To fully understand the different dynamics of systems that can occur over time, it is convenient to cover the arithmetic of complex numbers. For that reason we introduce this arithmetic in a seperate section. This analysis is used to introduce complex eigenvalues of a real square $n \times n$ matrix A. We show that each eigenvector has a generalized eigenspace. By choosing a basis for each of these generalized eigenspaces in an appropriate way we then obtain a basis for $\mathbb{R}^n$. With respect to this basis matrix A has the Jordan canonical structure. Since this Jordan canonical form has a diagonal structure it is a convenient way of analyzing the dynamics of a linear system (Chapter 3). Algebraic Riccati equations play a crucial role in this book. Therefore, we introduce and discuss a number of their elementary properties in the second part of Chapter 2. In this chapter we focus on the Riccati equation that is associated with the one-player linear quadratic control problem. We show that the solutions of this Riccati equation can be obtained by

determining the eigenstructure of this equation with the associated so-called Hamiltonian matrix. The so-called stabilizing solutions of Riccati equations play a crucial role later on. We show that the Riccati equation considered in Chapter 2 always has at most one stabilizing solution. Furthermore, we show that 'under some conditions' the Riccati equation has a stabilizing solution if and only if the associated Hamiltonian matrix has no eigenvalues on the imaginary axis.

Chapter 3 reviews some elementary theory on dynamical systems. The dynamics of systems over time are described in this book by sets of differential equations. Therefore the first question which should be answered is whether such equations always have a unique solution. In its full generality the answer to this question is negative. Therefore we review in Chapter 3 some elementary theory that provides us with sufficient conditions to conclude that a set of differential equations has a unique solution. To that end we first consider systems generated by a set of linear differential equations. Using the Jordan canonical form we show that such systems always have a unique solution.

For systems described by a set of nonlinear differential equations we recall some fundamental existence results from literature. A disadvantage of these theorems is that they are often useless if one is considering the existence of solutions for a dynamical system that is subject to control. This is because in most applications one would like to allow the control function to be a discontinuous function of time – a case which does not fit into the previous framework. For that reason the notion of a solution to a set of differential equations is extended. It turns out that this new definition of a solution is sufficient to study optimal control problems with discontinuous control functions.

Stability of dynamical systems plays an important role in convergence analyses of equilibrium strategies. For that reason we give in sections 3.3 and 3.4 an outline of how the behavior of, in particular planar, dynamical systems can be analyzed. Section 3.5 reviews some system theoretical concepts: controllability, stabilizability, observability and detectability. In section 3.6 we specify the standard linear quadratic framework that is used throughout this book. We show how a number of problems can be reformulated into this framework. Finally, we present in the last section of this chapter a number of examples of linear quadratic differential games which should help to motivate the student to study this book.

Chapter 4 deals with the subject of how to solve optimal control problems. The first section deals with the optimization of functions. The rest of the sections deal with dynamic optimization problems. As an introduction we derive the Euler–Lagrange conditions. Then, we prove Pontrjagin's maximum principle. Since the maximum principle only provides a set of necessary conditions which must be satisfied by the optimal solution, we also present some sufficient conditions under which one can conclude that a solution that satisfies the maximum principle conditions is indeed optimal.

Next, we prove the basic theorem of dynamic programming which gives us the optimal control of the problem, provided some conditions are met. It is shown how the maximum principle and dynamic programming are related.

Chapter 5 studies the regular linear quadratic control problem. The problem is called regular because we assume that every control effort is disliked by the control designer. We consider the indefinite problem setting, i.e. in our problem setting we do not make assumptions about preferences of the control designer with respect to the sign of deviations from the state variable from zero. The problem formulation allows for both

a control designer who likes some state variables becoming as large as possible and for a control designer who is keen on keeping them as small as possible. For both a finite- and infinite-planning horizon we derive necessary and sufficient conditions for the existence of a solution for this control problem. For the infinite-planning horizon setting this is done with the additional assumption that the closed-loop system must be stabilized by the chosen control. These existence conditions are phrased in terms of solvability conditions on Riccati equations. Moreover, conditions are provided under which the finite-planning horizon solution converges. We show that, generically, this solution will converge to the solution of the infinite-planning horizon problem.

Chapter 6 is the first chapter on differential games. It considers the case that players cooperate to achieve their goals in which case, in general, a curve of solutions results. Each of the solutions on this curve (the Pareto frontier) has the property that it cannot be improved by all the players simultaneously. We show for our linear quadratic setting how solutions of this Pareto frontier can be determined. Moreover, we show how the whole Pareto frontier can be calculated if all the individual players want to avoid the state variables deviating from zero. This can be done by solving a parameterized linear quadratic control problem.

Given this cooperative mode of play from the players the question arises as to how they will coordinate their actions or, to put it another way, which solution on the Pareto frontier will result. In section 6.2 we present a number of outcomes that may result. Different outcomes are obtained as a consequence of the fact that the sought solution satisfies different desired properties. For some of these outcomes we indicate how they can be calculated numerically.

Chapter 7 considers the case that the players do not cooperate to realize their goals. Furthermore, the basic assumption in this section is that the players have to formulate their actions as soon as the system starts to evolve and these actions cannot be changed once the system is running. Under these assumptions we look for control actions (Nash equilibrium actions) that are such that no player can improve his position by a unilateral deviation from such a set of actions. Given this problem setting we derive in section 7.2, for a finite-planning horizon, a both necessary and sufficient condition under which, for every initial state, there exists a Nash equilibrium. It turns out that if an equilibrium exists, it is unique. Moreover we show that during some time interval a Nash equilibrium exists if and only if some Riccati differential equation has a solution. A numerical algorithm is provided to calculate the unique Nash equilibrium actions.

For the infinite-planning horizon case things are more involved. In this case, if an equilibrium exists at all, it will in general not be unique. That is, in most cases there will exist an infinite number of Nash equilibrium actions. We show that if the equilibrium actions should permit a feedback synthesis, then the game has an equilibrium if and only if a set of coupled algebraic Riccati equations has a stabilizing solution; but, also in this case, there may exist an infinite number of equilibrium actions. A numerical algorithm is provided to calculate these equilibrium actions. A necessary and sufficient condition is given under which the game has a unique equilibrium. This equilibrium always permits a feedback synthesis.

Finally, we show that generically the finite-planning horizon equilibrium actions converge. If convergence takes place they usually converge to the actions implied by the infinite-planning horizon solution which stabilizes the system most. The chapter concludes with elaborating the scalar case and providing some examples from economics.

Chapter 8 also considers the non-cooperative mode of play. However, the basic assumption in this section is that the players know the exact state of the system at every point in time and, furthermore, they use linear functions of this state as a means of control to realize their goals. For a finite-planning horizon we show that this game has for every initial state a linear feedback Nash equilibrium if and only if a set of coupled Riccati differential equations has a symmetric set of solutions. Moreover, this equilibrium is unique and the equilibrium actions are a linear function of these Riccati solutions.

For the infinite-planning horizon case things are even more involved. In this case the game has for every initial state a Nash equilibrium if and only if a set of coupled algebraic Riccati equations has a set of symmetric solutions which, if they are simultaneously used to control the system, stabilize it. We elaborate the scalar case and show, in particular, that the number of equilibria may range from zero to $2^N - 1$. A computational algorithm is provided to calculate all Nash equilibrium actions. For the non-scalar case it is shown that there are games which have an infinite number of equilibrium actions.

Finally we show that in the two-player scalar game, if deviations of the state variable from zero are penalized, the solution of the finite-planning horizon game converges if the planning horizon converges. This solution always converges to a solution of the infinite-planning horizon game. However, different from the open-loop case, the converged solution now depends crucially on the scrap values used by both players.

Chapter 9 is the last chapter on non-cooperative games. It considers what effect uncertainty has on the equilibrium actions assuming that players are aware of the fact that they have to control a system characterized by dynamic quasi-equilibrium. That is, up to now we assumed that optimization takes place with no regard to possible deviations. It can safely be assumed, however, that agents follow a different strategy in reality. If an accurate model can be formed for all of the system, it will in general be complicated and difficult to handle. Moreover, it may be unwise to optimize on the basis of a model which is too detailed, in view of possible changes in dynamics that may take place in the course of time and that may be hard to predict. It makes more sense for agents to work on the basis of a relatively simple model and to look for strategies that are robust with respect to deviations between the model and reality.

We consider two approaches to model such situations. One is based on a stochastic approach. The other is based on the introduction of a malevolent deterministic disturbance input and the specification of how each player will cope with his aversion against this input.

We show that the equilibrium actions from Chapter 8 are also equilibrium actions for the stochastic counterparts of the games we study in this chapter. For the deterministic approach we see a more diverse pattern of consequences. Equilibrium actions may cease to exist for the adapted game, whereas opposite results are possible too. That is, a game which at first did not have an equilibrium may now have one or more equilibria. Sufficient existence conditions for such, so-called soft-constrained, Nash equilibria are provided. These conditions are formulated in terms of whether certain Riccati (in)equalities have an appropriate solution. For the scalar case, again, an algorithm is provided to calculate all soft-constrained Nash equilibria.

Finally, we show that the deterministic approach also facilitates the so-called linear exponential gaussian stochastic interpretation. That is, by considering a stochastic framework with gaussian white noise and players considering some exponential cost

function, the same equilibrium actions result. This result facilitates a stochastic interpretation of worst-case design and vice versa.

1.4 Notes and references

For the historical survey in this chapter we extensively used the MacTutor History of Mathematics Archive from the University of St. Andrews in Scotland (2003) and the outline of the history of game theory by Walker (2001). Furthermore, the paper by Breitner (2002) was used for a reconstruction of the early days of dynamic game theory.