THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Aspects of Waiting and Contracting in Game Theory

PETER HELGESSON

CHALMERS | UNIVERSITY OF GOTHENBURG

Division of Mathematics
Department of Mathematical Sciences
CHALMERS UNIVERSITY OF TECHNOLOGY
AND UNIVERSITY OF GOTHENBURG
Göteborg, Sweden 2015

Aspects of Waiting and Contracting in Game Theory
Peter Helgesson
ISBN 978-91-7597-208-4

Doktorsavhandlingar vid Chalmers tekniska högskola
Ny serie nr 3889
ISSN 0346-718X

Department of Mathematical Sciences
Chalmers University of Technology
and University of Gothenburg
412 96 Göteborg
Sweden
Phone: +46 (0)31-772 10 00

Printed in Göteborg, Sweden, 2015

Aspects of Waiting and Contracting in Game Theory
Peter Helgesson

Abstract

The topic of this thesis concerns two selected problem in game theory; *the N-player War of Attrition* and *the Principal-Agent problem*.

The War of Attrition is a well established game theoretic model that was first introduced in the 2-player case by John Maynard Smith. Although the original idea was to describe certain animal behaviour in for example territorial competition the model has also found interesting applications in economic theory. Following the results of Maynard Smith, John Haigh and Chris Cannings generalised the War of Attrition to two different models involving more than 2 players. We have chosen to call these generalizations the *dynamic*- and the *static model*.

In Paper I we study the asymptotic behavior of the N-player models as the number of players tend to infinity. By a thorough analysis of the dynamic model we find a connection to the more difficult static model in the infinite regime. This connection is then confirmed by approaching the limit of infinitely many players also in the static model. By using these limit results as a source of inspiration for the finite case, we manage to prove new results concerning existence and non-existence of an equilibrium strategy in the N-player static case.

Principal-Agent problems constitute a class of models in economics treating optimal incentives, often under moral hazard. A typical motivation is for instance optimal contract formation between employer and employee. In 1987 Bengt Holmström and Paul Milgrom wrote an influential paper on the subject, considering a discrete time model. This work has served as one of the main references for more recent studies on continuous time models.

Paper II and Paper III consider aspects of continuous time Principal-Agent problems by a stochastic maximum principle approach, thus not relying on the dynamic programming principle. In Paper II we introduce *the Hidden Action model* and *the Hidden Contract model* and characterize optimal contracts in these by following a scheme of sequential optimization. We also suggest a possible approach for solving a strong formulation of the Hidden Action model. Paper III generalizes the approach of Paper II to involve models having time inconsistent utility functions. By doing so we are able to consider Principal-Agent problems describing a risk averse behaviour, for instance in the sense of mean-variance.

Keywords: Game Theory, War of Attrition, ESS, N-Player, Principal-Agent Problem, Stochastic Maximum Principle, Pontryagin's Maximum Principle, Mean-Variance, Time Inconsistent Utility Functions

Preface

This thesis consists of the following papers.

- ▹ Peter Helgesson, Bernt Wennberg,
  "*The War of Attrition in the Limit of Infinitely Many Players*",
  in *Dyn. Games Appl.* **5** (2015), no. 1, 65–9-3.
- ▹ Boualem Djehiche, Peter Helgesson,
  "*The Principal-Agent Problem; A Stochastic Maximum Principle Approach* ",
  submitted to *SIAM J. Finan. Math.*
- ▹ Boualem Djehiche, Peter Helgesson,
  "*The Principal-Agent Problem With Time Inconsistent Utility Functions* ",
  preprint.

Acknowledgements

Five years have passed since the bright spring day when I was offered the opportunity to become a PhD-student in mathematics. The journey has taken me through a difficult and vertiginous terrain of high peaks, deep valleys and sudden impediments. It is therefore a pleasure to finally get to acknowledge all of you who have been encouraging me along the way and helping me reach my goal.

First and foremost I would like to thank my advisor Bernt Wennberg and my coadvisor Boualem Djehiche. Without your support this thesis would simply not have seen the light of day.

Thank you Bernt for accepting me as your student, for always being very friendly and for constantly pushing my ideas further. Especially, working with you on the project which resulted in my licentiate thesis was a time of true creativity, scientific curiosity and personal development.

Boualem, thank you for suggesting me to work on the Principal-Agent problem and for your keen hospitality during my many visits to KTH. Your engagement in my progress, giving me a call every now and then, if only just for a quick chat, has meant a lot.

Next I would like to express my deepest gratitude to all of my present and former colleagues at the department of mathematics for contributing to such a stimulating and friendly atmosphere. Some of you have been close friends of mine even since before my time here. Therefore, a very special thank you to:

Dawan, for sharing the office with me through these years and for being one of my best friends. Having you around, wether it has been during the long hours of traveling to conferences, or in the tight moments of racing each other on the bikes, has always been a pleasure. I cannot thank you enough.

Hossein, for your great company, both at work and in private, and for our many good discussions on literature, teaching and life.

Magnus Ö, for always making me push myself a bit further in all of our many common interests and for installing my router.

Oskar, for your sociable personality and for diligently having carried the role as commander in chief for many of our Friday after-works.

Richard, for being a good friend and a true role model, both as a mathematician, but foremost, also as a person.

Magnus R, for being the first to encourage me studying mathematics outside of the Chalmers program and for sharing my passion for live music.

In fact, I have made many good friends here at the department, all having contributed to make my time run all too fast: Matteo, Claes, Ivar, Malin, Jakob, Timo, Adam, Jonas, Anna, Henrike, Philip, Elizabeth, David, Emil, Ragnar, Aron and many others. Thank you all!

During my time as an undergraduate student I was fortunate enough to meet several great teachers. I would especially like to thank Mats Andersson, Christer Borell, Jana Madjarova, Christian Söderberg and Lars Vretare for being among those. Your passion for the subject and enthusiasm in teaching it to others has been a constant source of inspiration and is something that I will never forget.

Another important element of my time here, which has been a most valuable weakly activity for keeping me sane, is the "lunchinnebandy". I would like to express a sincere thank you to Jonathan, Gustav, Henrik and Patrik for always showing up and giving a hard time.

Finally, an enormous thank you to my wonderful family for your endless love and support; my parents Tomas and Carina, my siblings Maria and Andreas, my grandparents and Linda. Having you around really makes my life complete. I love you.

Peter Helgesson
Göteborg, May 2015

Contents

Part I

INTRODUCTION

Introduction

Game theory, when defined in its broadest sense, could be thought of as a collection of models formulated to study situations of conflict and co-operation. By analysing these models game theorists try to find answers to questions concerning best actions for individual decision making, but also to what mechanisms that could underlie social behaviour. This introduction aims at giving a concise survey of some fundamental parts of the subject and present the most central results. We begin by a quick journey in history.

1. Brief History

When writing a text on the history of a specific scientific subject it is inevitable not to miss out or to exclude interesting pieces of it, no matter how hard you try. Game theory is not an exception. It is even difficult to say for sure when and where it began since, in a wide perspective, strategic thinking has always been around. The foundation for what game theory is today therefor rests in hundreds, and even thousands, of years of history. The first contribution to mathematical game theory is, however, often credited to the French mathematician Antoine Augustin Cournot (1801-1877) for his book *Recherches sur les principes mathématiques de la théorie des richesses* from 1838. Cournot introduces mathematical tools (functions, probabilities, etc.) in the context of economic analysis and, most importantly, he constructs a theory of oligopolistic firms and analyses oligopolistic competition. Thirty years after its publication these ideas were to have a strong influence on, what was to become, modern national economy.

One of the first contributions to what we recognise as pure classical game theory today was made in a series of papers and notes by Émile Borel (1871-1956) during the period 1921-1927. Borel studied finite symmetric 2-player games at an abstract and more general level, without having any particular application in mind. In his work he introduced the concept of "*méthode de jeu*" (method of game, strategy) which he used to pose the fundamental question of whether it was possible to determine a "*méthode de jeu meilleure*" (best method of game) or not, even though it was not properly defined what a "*méthode de jeu meilleure*" actually was. Strange as it is, the work of Borel was not recognised until years after the breakthrough by von Neumann that was to come. In 1953, while translating Borel's work to English, the French mathematician Maurice Fréchet is quoted saying "... in reading these notes of Borel's I discovered that in this domain [game theory], as in so many others, Borel had been an initiator." (see [**F4**]).

The first huge impact in the mathematical theory of games came in year 1928 through the works of John von Neumann in [**vN**]; *Zur Theorie der Gesells- chaftsspiele*. In this paper he gives a complete proof of the classical "minimax theorem" for 2-player zero-sum games, basically saying that Borel's "best method of game" indeed always exists in the zero-sum case. This result is probably the most influential in the history of game theory and it would not be an overstatement to claim that the subject was born in 1928 through the paper of von Neumann. Apart from the minimax theorem von Neumann was the first to clearly explain the passing from extensive-form games to the more useful notion of normal-form games. The normal-form was to become of great importance not only for game theory, but also for the shaping of modern economic theory.

In their famous book *Theory of games and economic behavior* from 1944 von Neumann and Morgenstern present the state of the art theory available at the time, including cooperative games with definitions of *TU-games* and the solution concept of *stable sets*. A short survey of von Neumann's contributions to game theory is given in e.g. [**F2**].

During the 30's and 40's much of the research done in game theory was focused on cooperative games in which players are engaged into coalitions. Even though this analysis was (and still is) both interesting and important it was somewhat leading away from other interesting questions of games where negotiation and individual decision making based on personal information is present. This direction gained momentum in the 1950's thanks to the contributions of John F. Nash to non-cooperative game theory and his famous equilibrium theorem. In 1951 Nash published the paper *Non-Cooperative Games* in which he defined and proved the general existence of a "best play" *equilibrium strategy*, or *Nash-equilibrium**, valid in all finite

*The concept was actually formulated already by Cournot, but in a much less general setting. Some scholars have suggested the name Cournot-Nash-equilibrium, or even Cournot-equilibrium, but most people agree that the depth of the definition is due to Nash.

normal-form games. This result was the ultimate generalization of von Neumann's minimax theorem, but it was also the igniting spark for future research in non-cooperative games (even though the impact of the result initially spread slowly).

Among many significant contributions by Nash (that range even outside game theory) he gave strong arguments as to why any theory of games should be reducible to equilibrium analysis of a normal-form game. He also gave a beautiful (axiomatic) argument to solve the so called two-person *bargaining problem*. For a good further reading of the importance that Nash's work has had in game theory and economics the text [**M**] by Myerson is recommended.

By the mid 50's and onwards game theory had become a well established area within the mathematical community and persons like John Harsanyi, Reinhard Selten, Robert Aumann and Lloyd Shapley, only to mention a few, made important contributions to the subject.

Considering applications the spectrum was initially rather narrow, mainly concentrated in economic theory. It would take until the 1970's for this to change by the works of two mathematical biologists, John Maynard Smith and George R. Price, and their definition of *evolutionary stable strategy*, or *ESS*, in "*The logic of animal conflict*" (see [**MSP**]). What Maynard Smith and Price realised was that ideas from game theory could be used to formulate a related dynamical theory, potentially useful for describing population dynamics. The big difference was that the players in this model were not assumed to act in a rational manner. This was the starting point of what today is known as *evolutionary game theory* which, apart from its original intention of being a tool in theoretical biology, has found applications also in economy, social science and philosophy.

Ever since its first major breakthrough in 1928 game theory has continued to expand both in terms of applications and theoretical development. Today there is a wide range of ongoing research of all kinds in subjects from classical- and evolutionary game theory to stochastic- and differential games. As late as in 2006 a new theory (related to differential games) was developed independently in works by P. L. Lions, J. M. Lasry and P. E. Caines, M. Huang, R. Malhamé, called *mean field game theory*, or just *MFG*. MFG has attracted lots of attention by opening doors to many potential applications, but the theory is juvenile and still somewhat under construction (the book [**BFY**] by Bensoussan, Frehse and Yam is one of few references to the subject apart from research papers).

2. Extensive-Form and Normal-Form games

Games of various kinds have been present in our society for ages in connection to economy and gambling, in strategic decision making and conflict scenarios, but quite often also just for the sake of fun. One of the most prominent of those games is undoubtedly chess, which has diverted mankind for centuries. The rules are so simple that one can learn about how to play only in a few minutes yet, not even a lifetime of practise is enough to fully master the complexity of the game. The difficulty of playing chess originates from the enormous number of moves a player can make during a play. Indeed, if we were to reduce the number of pieces for each player from 16 to 8 the game would become rather poor. Apart from the matter of "size", theoretically speaking, chess is actually very simple. Its general structure is common for a wide range of other 2-player games such as for instance Othello, Nim, Go etc., which are all typical examples of games that can be represented in a so called *extensive-form*. Loosely speaking an extensive-form 2-player game is a finite directed tree where each node represents a player in a specific position during a game. There are three different types of nodes; one having outgoing edges but no incoming called the *root*, nodes having both outgoing and incoming edges are *intermediate*, and nodes having an incoming edge but no outgoing are called *terminal*. An edge connecting two nodes represent a move that the player at the first node can use to get to the second node. The root of the tree is the starting point of the game. In chess for instance the white player is in the root from which there are 20 different moves leading to an intermediate node for the black player. The game is played via intermediate nodes until it finally reaches a terminal node where the game ends (checkmate or draw).

The outcome of a play is measured by means of a *payoff function* which assigns a 2-vector to each of the terminal nodes. The elements in this vector represent the payoffs given to each of the players respectively. If we again use chess as a game of reference we could assign values to each possible outcome as "win = 1", "draw = 0" and "loose = -1", and hence get $\{(1,-1),(-1,1),(0,0)\}$ as the set of possible values of the payoff function.

At this point we have all the ingredients required to describe the simplest kinds of 2-player games in extensive-form, namely; a game tree, on which the game is being played, and a pay-off function, measuring the outcome of a play. This is indeed a good start, but in order to find a complete theory of games it is obvious that we should be looking for a more general description.

Take for instance the game of poker in 4 players. Then we face two new features that our simple description can not yet meet. Firstly, the number of players is greater than two and, secondly, the players are unaware of how their opponents are playing. One could also think of games having "chance moves", i.e. random moves that does not connect to any particular player. The following definition of a game in extensive-form covers all of the above features and can be found in [**O**]:

Definition 1. By an n-player game in extensive form we mean

(1) a topological tree Γ with a distinguished node A called the *starting point*, or the *root*, of Γ
(2) a function, called the *pay-off function*, which assigns an n-vector to each terminal node of Γ
(3) a partition of the intermediate nodes of Γ into $n+1$ sets $S_0, S_1, ..., S_n$, called the player sets
(4) a probability distribution, defined at each node of S_0, among the immediate followers of this node
(5) for each $i = 1, ..., n$ a subpartition of S_i into subsets S_i^j, called *information sets*, such that two nodes in the same information set have the same number of immediate followers and no node can follow another node in the same information set
(6) for each S_i^j there is an index set I_i^j together with a 1-1 mapping of I_i^j onto the set of immediate followers of each node in S_i^j.

As mentioned earlier condition (1) and (2) suffice to describe the simplest games in extensive form like e.g. chess. Condition (3) sets the stage for the n-player generalization where S_i, $i \neq 0$, should be thought of as the collection of nodes in Γ from which player i makes a move. The set S_0 differs from $S_1, ..., S_n$ in that it contains nodes from which the game proceeds at random (without any player making a move) to an immediately following node, i.e. S_0 is the collection of chance nodes. Conditions (5) and (6) open up the possibility of having a "lack of knowledge" in the game. For $i \neq 0$ one should think of the nodes in S_i^j as different positions in the play of player i that are indistinguishable. In poker for instance every move (that is not terminal) is followed by a chance move (drawing a card) and the only information available to a given player i is what cards he has at the moment and what cards he has decided to discard. No information of the opponents hands is available so, for a fixed hand, all the possible nodes of player i in a round j for which the same cards have been discarded by i (in any order) is indistinguishable to him. Thus there is a natural partition of S_i, for each $i = 1, ..., n$, where each S_i^j in the partition contains several nodes of Γ. A game in which $|S_i^j| = 1$ for all i and j, i.e. all nodes are always distinguishable, is said to have *perfect information*. Chess is a typical example of a game under perfect information.

Using the terminology of Definition 1 we are now ready to introduce the fundamental concept of *strategy*.

Definition 2. Let Γ be an n-player game in extensive form and let S_i^j, for $1 \leq i \leq n$, be the information sets of a player i. A *(pure) strategy* of player i is defined as a function σ_i from each S_i^j to any of the edges which follow a representative node of S_i^j. The set of all strategies available to player i is denoted by Σ_i. An element in the product space $\Sigma_1 \times \Sigma_2 \times ... \times \Sigma_n$ is called a *profile* of strategies.

The above definition captures the intuitive idea of what a strategy "should be", namely; a strategy is a complete plan of how to play in any given situation. There is, however, also a drawback with Definition 2 since it assumes the player to have decided about how to play even before the game has started. One may of course argue that this is unreasonable in many practical situations. Again chess is a good example since one would not be able to make up a plan for what to do in more than a few moves ahead. Indeed this is a practical limitation that we will have to overlook. It will not make the theory less interesting.

The introduction of strategies is of course of fundamental importance since they represent the basic elements of what game theory means to study. We want find out wether there is a best way of choosing strategies or not, and how to define this property. From the point of view of a player this would be to pick a strategy which maximises the personal payoff. Given that the opponents play according to some profile of strategies one would pick a strategy so that the final position gets as good as possible. For this discussion it is time to introduce some further notation. We can make the notion of payoff-function in an n -player game precise by declaring it as a function $\mathscr{J} : \Sigma_1 \times \Sigma_2 \times ... \times \Sigma_n \to \mathbb{R}^n$ where

$$\mathscr{J}(\sigma_1, ..., \sigma_n) = (\mathscr{J}_1(\sigma_1, ..., \sigma_n), ..., \mathscr{J}_n(\sigma_1, ..., \sigma_n)),$$

and $\mathscr{J}_i$ is the payoff-function of player i. Note that since Γ may consist of chance moves one should, in general, interpret $\mathscr{J}$ as an expectation.

In many situations, both in theory and in practise and given that every player have chosen a strategy, we are only interested in the values of each individual payoff-function. In principle, given the product space $\Sigma_1 \times \Sigma_2 \times \ldots \times \Sigma_n$ and the payoff-function $\mathscr{J}$ we have a characterisation of the game which in many ways is sufficient for our needs. Describing a game by the payoff-function is commonly known as a *normal-form* representation. Every game in extensive form can canonically be represented in normal form, but not in the converse order. The reason for this being that a game in extensive form contains information of the game tree Γ which is lost in the normal form representation. If the number of strategies in each Σ_i is finite (which we have assumed) the normal-form representation is simply given by an n-dimensional array of n-vectors. In the case of a 2-player game this array reduces to a bimatrix.

Example 1. One of the most famous examples of models in game theory is undoubtedly *The Prisoner's Dilemma*, first introduced in 1950 by Merrill Flood and Melvin Dresher at the RAND Corporation. The story goes as follows: Two members of a criminal gang are arrested by the police and imprisoned in two different rooms without being able to communicate with each other. Both of them can act in one of the following ways: either choose to *cooperate*, i.e. keep quite to the police during the interrogation, or else *defect* and choose to testify against the other prisoner. We denote cooperation by C and defection by D. If the prisoners cooperate they will both go to jail for two years and if they defect they will get three years in prison. The catch of the game is that if one of the prisoners choose to cooperate while the other one defects, then the latter will be released while the cooperative prisoner will get four years behind the bars. Denoting the prisoners by Player 1 and Player 2 this game can easily be represented in both extensive- and normal form as in Fig. 1.
The dashed line in the extensive form representation indicates that the nodes belong to the same information set. For Player 2 these nodes are indistinguishable. Such information is not included in the normal-form representation.

We are now ready to introduce the notion of Nash-equilibrium.

Definition 3. Given an n-player game Γ we say that an n-tuple of strategies $(\sigma_1^*, \ldots, \sigma_n^*) \in \Sigma_1 \times \Sigma_2 \times \ldots \times \Sigma_n$ is a *(pure) Nash-equilibrium* if and only if for

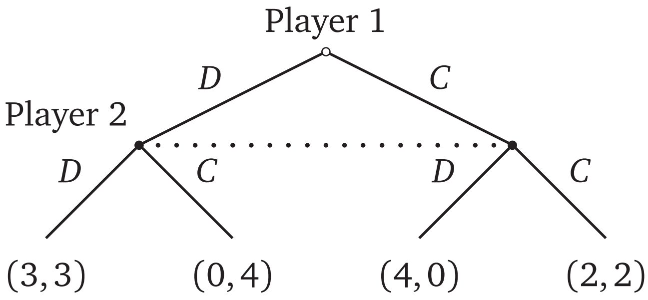

	C	D
C	2,2	4,0
D	0,4	3,3

FIGURE 1. The Prisoner's Dilemma represented in extensive- and normal form.

any $i = 1, \ldots, n$ and $\sigma_i \in \Sigma_i$,

$$\mathscr{J}_i(\sigma_1^*, \ldots, \sigma_n^*) \geq \mathscr{J}_i(\sigma_1^*, \ldots, \sigma_{i-1}^*, \sigma_i, \sigma_{i+1}^*, \ldots, \sigma_n^*).$$

Nash-equilibrium is one of the most celebrated definitions in non cooperative game theory and serves as the major solution concept. The intuitive meaning is clear. Given that all players in the game play according to their Nash-equilibrium strategy, none of them can get a higher (expected) payoff by switching to another strategy. Despite its great importance, Nash-equilibrium is far from being the only solution concept in the litterateur. Other concepts are for instance *subgame perfect equilibrium* and *evolutionary stable strategy*.

At this point an important question naturally rises; given a finite game Γ, does there always exist a Nash-equilibrium? A moment's thought will revile this to be false. Consider for instance the 2-player normal-form game:

$$\begin{pmatrix} (1,-1) & (0,0) \\ (0,0) & (1,-1) \end{pmatrix} \tag{2.1}$$

where, at each entry, either Player 1 or Player 2 can do better by changing to another strategy. There is, however, more to be said about this matter as we will see shortly.

We conclude this section with the important notion of symmetric games.

Definition 4. Let Γ be an n-player game given in normal form. We say that Γ is *symmetric* if and only if for every $i = 1, \ldots, n$ and permutation π it holds that

$$\mathscr{J}_i(\sigma_1, \sigma_2, \ldots, \sigma_n) = \mathscr{J}_{\pi(i)}(\sigma_{\pi(1)}, \sigma_{\pi(2)}, \ldots, \sigma_{\pi(n)}),$$

for all $(\sigma_1, \sigma_2, \ldots, \sigma_n) \in \Sigma_1 \times \Sigma_2 \times \ldots \times \Sigma_n$.

3. Zero-Sum Games

The simplest possible games to study are the so called *zero-sum games*. They are characterised by the fact that the sum of the elements of the payoff-vector at any terminal node always equals to zero. Thus, in a zero-sum game the winnings of one player has to be paid, in some way or another, by the other players. If only 2-player zero-sum games are considered, things simplify even further since the elements of the payoff-vector then have to be additive inverses of each other. In that case it suffices to give the payoff of one player, say Player 1. The normal-form bimatrix representation can therefore be reduced to a matrix representation. For this reason zero-sum games are sometimes also referred to as *matrix games*. The game in (2.1) is a zero-sum game and in the reduced notation of the first player's payoff we get the game matrix

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}. \tag{3.1}$$

The easy matrix representation makes the 2-player zero-sum games especially tractable for closer analysis. For instance to answer questions regarding the existence of solution strategies. We know that (3.1) does not admit a pure Nash-equilibrium but maybe we can quantify those matrix games that have? Consider a general zero-sum game with a pay-off function $\mathscr{J}$ corresponding to a payoff-matrix $\mathbf{A} = (a_{ij}) \in \mathbb{R}^{n\times m}$. The strategy sets of Player 1 and Player 2 are Σ_1 and Σ_2 respectively and both are assumed to be finite. The pair $(\sigma_1^*, \sigma_2^*) \in \Sigma_1 \times \Sigma_2$ is a pure Nash-equilibrium if and only if both $a_{i^*j^*} = \max_i a_{ij^*}$ and $a_{i^*j^*} = \min_j a_{i^*j}$, where $\mathscr{J}(\sigma_1^*, \sigma_2^*) = a_{i^*j^*}$. Such an element, if it exists, is called a *saddle point* of $\mathbf{A}$. If $\mathbf{A}$ lacks saddle points the game lacks pure Nash-equilibria. What would happen if we were playing such a game? The goal of Player 1 is to win as much as possible while minimising the risk of loosing too much. Thus, in each row two elements are of interest; the greatest (maximal gain) and the least (maximal loss). A rational choice of strategy for Player 1 would be to pick a strategy corresponding to the row in $\mathbf{A}$ in which the least possible win is maximised. An analogue argument also holds for Player 2 who preferably would choose to play a strategy corresponding to the column in which the greatest loss is minimised. We define

$$\underline{v} := \max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} \mathscr{J}(\sigma_1, \sigma_2)$$
$$\overline{v} := \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} \mathscr{J}(\sigma_1, \sigma_2)$$

and call $\underline{v}$ the *gain-floor* and $\overline{v}$ the *loss-ceiling*. By construction we have the inequality

$$\underline{v} \leq \overline{v}, \tag{3.2}$$

and if "=" in (3.2) we say that the game has *value* $v = \underline{v} = \overline{v}$. Hence, if a zero-sum game has a value the payoff-matrix has a saddle point and there exists a pure Nash-equilibrium. If not, then the gain-floor and the loss-ceiling only represent the best possible win and loss each of the players can hope for. However, there is a way of playing games without saddle points so that both players gain from it. That is the concept of *mixed strategies*, first introduced by Émile Borel:

Definition 5. Let Γ be an n-player game in normal form with strategy spaces $\Sigma_1, \Sigma_2, ..., \Sigma_n$ and payoff-function $\mathscr{J}$. We say that μ_i is a *mixed strategy* of player i if $\mu_i \in \mathscr{M}_1(\Sigma_i)$, where $\mathscr{M}_1(\Sigma_i)$ is the set of probability measures over the space Σ_i.

Note that in mixed strategies the payoff-function $\mathscr{J}$ turns into an expected payoff-function given by

$$\int_{\Sigma_1 \times \Sigma_2} \mathscr{J}(\sigma_1, \sigma_2) \mu_1(d\sigma_1) \mu_2(d\sigma_2),$$

for $(\mu_1, \mu_2) \in \mathscr{M}_1(\Sigma_1) \times \mathscr{M}_1(\Sigma_2)$. For simplicity we stick to the same notation as for the ordinary payoff-function.

To play a mixed strategy is indeed a bit odd. Practically it means that instead of using rational reasoning to find a good strategy one would draw a strategy at random according to some probability distribution. By the inclusion $\Sigma_i \subset \mathscr{M}_1(\Sigma_i)$, there is reason to believe that the mixed strategies enable us to find values in a wider class of games. The following result is due to von Neumann in [**vN**] (1928) and can be considered as the fundamental theorem of game theory.

Theorem 3.1 (The minimax theorem). *Let Γ be a 2-player zero-sum game with finite strategy spaces Σ_1 and Σ_2 and payoff-function $\mathscr{J}$. Then there exists at least one Nash-equilibrium in mixed strategies $(\mu_1, \mu_2) \in \mathscr{M}_1(\Sigma_1) \times \mathscr{M}_1(\Sigma_2)$*

for which the game has value, i.e.

$$\max_{\mu_1 \in \mathcal{M}_1(\Sigma_1)} \min_{\mu_2 \in \mathcal{M}_1(\Sigma_2)} \mathcal{J}(\mu_1, \mu_2) = \min_{\mu_2 \in \mathcal{M}_1(\Sigma_2)} \max_{\mu_1 \in \mathcal{M}_1(\Sigma_1)} \mathcal{J}(\mu_1, \mu_2).$$

The minimax theorem was the first major breakthrough in what was to become the theory of games and it has later been generalized by several authors. John von Neumann himself was quoted as saying "As far as I can see, there could be no theory of games ... without that theorem ... I thought there was nothing worth publishing until the minimax theorem was proved" (see [**C2**]).

Since we have assumed the number of strategies available to each player to be finite (n say) any mixed strategy may be represented by a vector $\mathbf{x} \in \mathbb{R}^n$ such that all $x_i > 0$ and $\sum_{i=0}^{n} x_i = 1$. Each element x_i is the probability of getting the pure strategy indexed by i when playing $\mathbf{x}$. The payoff-function can be written as

$$\mathcal{J}(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \mathbf{A} \mathbf{y}, \qquad \mathbf{x} \in \mathbb{R}^n, \mathbf{y} \in \mathbb{R}^n.$$

The minimax theorem is a pure existence result, but there is also another interesting result saying that equilibrium strategies in symmetric 2-player zero-sum games can be derived as solutions to a certain ODE-system. The following theorem can be found in [**F1**] and is due to von Neumann:

Theorem 3.2. *Let $\mathbf{A} \in \mathbb{R}^{n\times m}$ be the payoff-matrix in a 2-player zero-sum game and define the functions $u_i(\mathbf{y}) = \mathbf{e}_i^T \mathbf{A}\mathbf{y}$ (where $\mathbf{e}_i$ denotes the i:th normalized column vector from the standard basis), for $i = 1, 2, ..., n$, $\phi(a) = \max(0, a)$ and $\Phi(\mathbf{y}) = \sum_{i=1}^{n} \phi(u_i(\mathbf{y}))$. For any mixed strategy $\mathbf{y}^0$ of Player 2 consider the following problem:*

$$\begin{cases} y_j'(t) = \phi\left(u_j(\mathbf{y}(t))\right) - \Phi\left(\mathbf{y}(t)\right) y_j(t) \\ y_j(0) = y_j^0 \end{cases}$$

Then, for any positive monotone sequence $\{t_k\}$ growing to infinity any limit point of $\{\mathbf{y}(t_k)\}$ is an equilibrium strategy of Player 2 and, furthermore, there is a constant C such that $\mathbf{e}_i^T \mathbf{A}\mathbf{y} \leq \sqrt{n}/(C + t_k)$.

It should be mentioned that finding the equilibrium solution by means of ordinary differential equations is not very efficient. Much faster solution algorithms have been developed using methods from linear programming.

The literature of 2-player zero-sum games is huge, both from a theoretical- and an applied point of view, and the topic serves as the foundation for what is called *classical game theory*. Apart from the 2-player setup there are also results available for n-player zero-sum games, but not as extensive. A substantial part of the classical text *Theory of games and economic behavior* by von Neumann and Morgenstern (see [**vNM**]) is, however, devoted to n-player problems.

4. Nash's Theorem

Following the footsteps of von Neumann the young Princeton mathematician John F. Nash was to give game theory its next major breakthrough by a beautiful generalisation of the minimax theorem. Based on Nash's work in his doctoral thesis the following result can be found in [**N**]:

Theorem 4.1. *Every normal-form n-player game with finitely many strategies has at least one Nash-equilibrium in mixed strategies.*

The proof of Nash's theorem builds upon an elegant fixed point argument, preferably using Brouwer's or Kakutani's fixed point theorem. The statement is true even for more general normal-form games having infinite strategy sets. Ever since its publication in 1951 the result has generated much attention and has by now created a whole avenue of interesting research in n-player non-cooperative games.

5. Evolutionary Game Theory

Classical non-cooperative game theory typically deals with problems concerning equilibrium analysis. Questions like; Is there a Nash-equilibrium? Is it unique? What is the expected payoff and what is the risk when playing a Nash-equilibrium? etc. are fundamental. From a game theorist's point of view these questions are of course very natural to ask, but are they equally natural for the economist or the evolutionary biologist when trying to understand actual social behaviour? Do agents really play Nash-equilibrium in a given situation and if they do, which one do they chose if several exist? These are some of the basic problems of interest in what is called *evolutionary game theory*. In contrast to classical theory, evolutionary games do not assume players to act in a rational manner. In every day life we are all being exposed to new situations which, in principle, could be analysed game theoretically. However, most often we do not think of such situations in strategic terms. It is therefore not realistic to assume our behavior to be reflected by rationality, since what we really use are often just simple strategies, imitation, experience and rules of thumb. The best one can hope for is that rational behaviour, as described by Nash-equilibrium, evolves over time. In evolutionary game theory the basic setup is a large population of players who repeatedly engage in strategic interaction. Change of behaviours in these populations is driven on an individual level by features such as imitation of more successful behaviours. For a thorough discussion of the passing from classical- to evolutionary games we recommend [**M**].

We are now going to present the basic model. Consider a finite population of N players and a normal form game admitting a finite set of (pure) strategies $\Sigma = \{\sigma_1, \sigma_2, ..., \sigma_n\}$. For simplicity we are going to assume all the individual strategy sets to be the same and equal to Σ. A *population state* x is a point in the n-dimensional unit simplex X and describes the proportion of each strategy being used within the population, i.e. x_i is the proportion of players using strategy σ_i. Note that for $N < \infty$ the values of x are in the grid $\mathscr{X}_N := X \cap \frac{1}{N}\mathbb{Z}^n$, embedded in X. A *population game* is a continuous vector-valued payoff-function $F : X \to \mathbb{R}^n$ for which each $F_i(x)$ is the (expected) payoff to a σ_i-player given a population state x. In this context a Nash-equilibrium is a population state x^* which fulfills the implication

$$x_i^* > 0 \quad \Rightarrow \quad i \in \operatorname*{argmax}_{1 \le j \le n} F_j(x^*),$$

for all $1 \le i \le N$ (see e.g. [**S2**]). It is easy to show that the above definition coincides with the classical Nash-equilibrium if the population game is a finite symmetric normal-form game.

As we mentioned earlier the goal of evolutionary game theory is to understand the possible mechanisms driving strategic behaviour within a population. How and why do players switch from one strategy to another and how do we model it? To deal with this we introduce *revision protocols*. Formally, a revision protocol is a map $\rho : \mathbb{R}^n \times X \to \mathbb{R}_+^{n\times n}$ taking payoff vectors π (π_i is the expected payoff to a σ_i-player given x) and population states x as arguments and returns a square matrix having positive elements. The idea is to consider a population of N individuals, each of them being equipped identical "alarm clocks". The time durations between two consecutive rings of a clock are independent and exponentially distributed with exponential parameter R. As soon a a clock rings the player carrying it gets a chance to switch to another strategy. The switching process is random and related to ρ in a way so that the probability of changing from strategy σ_i to σ_j, $i \neq j$, is given by ρ_{ij}/R. When a switch occurs at a time t the population state vector also reacts by jumping to a neighbouring point in $\mathscr{X}_N$. This jump-process can be described as a continuous time Markov chain which we denote by

$X_N(t)$. The aim is to study the time evolution of $X_N(t)$ and especially its asymptotic behaviour when $t \to \infty$. At this point there are two possible routes to take; either we study $X_N(t)$ directly, the so called *stochastic dynamics*, or we study the expectation $x_t := \mathbb{E}[X_N(t)]$, the *mean dynamics*. In this text we are going to focus on the latter in the case of population games.

Given a revision protocol ρ and a population game F there should be an equation for x_t. To find it we consider the expected differential of $X_N(t)$ over a small time interval $[t, t+dt]$. In dt units of time the expected number of revision opportunities for each player is Rdt. Thus, given a population state x at time t, there will be on average Nx_iRdt revision opportunities within the group of σ_i-players. Therefore, since the probability of changing strategy is ρ_{ij}/R, we expect $Nx_i\rho_{ij}dt$ of the players in this group to switch to σ_j in the time interval $[t, t+dt]$. Adding up the expected number of immigrants and emigrants to and from strategy σ_i one finds

$$N dx_i = N\left(\sum_{j=1, j\neq i}^{j=n} x_j\rho_{ji}(F(x), x) - x_i \sum_{j=1, j\neq i}^{j=n} \rho_{ij}(F(x), x)\right) dt$$

which yields following differential equation:

$$\dot{x}_i = \sum_{j=1, j\neq i}^{j=n} x_j\rho_{ji}(F(x), x) - x_i \sum_{j=1, j\neq i}^{j=n} \rho_{ij}(F(x), x) =: V^F(x), \tag{5.1}$$

forming a system of n ordinary differential equations called *the mean dynamic*. A population state x such that $V^F(x) = 0$ is called a *stationary point*.

So far nothing particular has been said about the revision protocol and its properties. The explicit form of ρ depends on what problem one would like to study and on what application one has in mind. It must therefore be constructed on a case to case basis. There is a handful of models in the literature of certain interest. In the context of evolutionary biology the most common model by far is the so called *replicator dynamics* (first introduced in [**TJ**] by Taylor and Jonker) which is generated from (5.1) by choosing $\rho_{ij}(\pi, x) = x_j[\pi_j - \pi_i]_+$. The basic idea behind this choice is simple; the probability of switching from strategy i to strategy j should be proportional to the proportion of players using σ_j (imitation) and to the advantage in payoff of playing σ_j instead of σ_i (payoff-advantage). If, given x, σ_j-players do worse than σ_i-player the probability of a switch is zero. Inserting this protocol in (5.1) we get the replicator dynamics:

$$\dot{x}_i = x_i \left(F_i(x) - \sum_{j=1}^{n} x_j F_j(x) \right). \tag{5.2}$$

Note that the form of (5.2) makes it impossible for strategies that are not present in the initial population to emerge later on.

In the case of a linear population game, i.e. $F(x) = \mathbf{A}x$ for some matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, the replicator equation can be written

$$\dot{x}_i = x_i \left((\mathbf{A}x_t)_i - x_t^T \mathbf{A} x_t \right). \tag{5.3}$$

The system created from (5.3) satisfy the following properties (see [**HS**]):

(1) if x is a Nash-equilibrium, then $V^F(x) = 0$
(2) if x is a strict Nash-equilibrium, then it is asymptotically stable
(3) if $V^F(x) = 0$ and x is the limit of an orbit in the interior of the simplex X as $t \to \infty$, then x is a Nash-equilibrium
(4) if $V^F(x) = 0$ and x is stable[†], then it is a Nash-equilibrium.

Note, however, that the converse implications of (1) - (4) are all false. The replicator dynamics does therefore not guarantee convergence of solutions to a Nash-equilibrium. The following is a simple example of such a situation.

Example 2 (Rock-Paper-Scissor). In the classic game of rock-paper-scissor there are obviously three pure strategies to chose from. We identify each of them by the unit vectors in $\mathbb{R}^3$: e_1, e_2 and e_3. The game is characterised by the fact that e_1 wins against e_2, e_2 wins against e_3 and e_3 wins against e_1. In the general setup the individual payoff of playing e_i against e_j can be written like $\mathscr{J}(e_i, e_j) = e_i^T \mathbf{A} e_j$ where

$$\mathbf{A} = \begin{pmatrix} 0 & -a_2 & b_3 \\ b_1 & 0 & -a_3 \\ -a_1 & b_2 & 0 \end{pmatrix}$$

for any $a_1, a_2, a_3, b_1, b_2, b_3 > 0$. This game has a unique Nash-equilibrium x^* in the interior of the unit simplex X (if for instance $a_1 = a_2 = a_3$ and $b_1 = b_2 = b_3$ it is $(1/3, 1/3, 1/3)$) which is asymptotically stable if and only if $\det \mathbf{A} > 0$. In the case $\det \mathbf{A} < 0$ the solutions of the replicator

[†]A point $z \in X$ is *stable* if for every neighbourhood U of z there exists another neighbourhood V of z such that if $x \in V$ then $x(t) \in U$ for all $t \geq 0$. Moreover, a state is *asymptotically stable* if it is a stable attractor.

dynamics, when starting at any interior state (not equal to x^*), will spiral to the boundary of X and never settle at the equilibrium state. This is illustrated in Fig. 2.

The inability of the replicator dynamics to ensure convergence to a Nash-equilibrium should not be considered as a flaw. It is merely an indication that it takes more than imitation of success to reach game theoretic rationality, which is an interesting observation in itself.

The replicator equation (5.2) is said to be *permanent* if there is a compact set $K \subset \text{int } X$ such that for all $x_0 \in \text{int } X$ there is a $T > 0$ such that for all $t > T$ one has $x_t \in K$. For such a dynamics (with F linear) we have the following (see [**HS**]):

Theorem 5.1. *If (5.3) is permanent, then there exists a unique stationary point $z \in int\, X$. The time averages along each internal orbit converge to z:*

$$\frac{1}{T}\int_0^T x_i(t)dt \xrightarrow{T \to \infty} z_i, \quad \textit{for } i = 1, 2, ..., n.$$

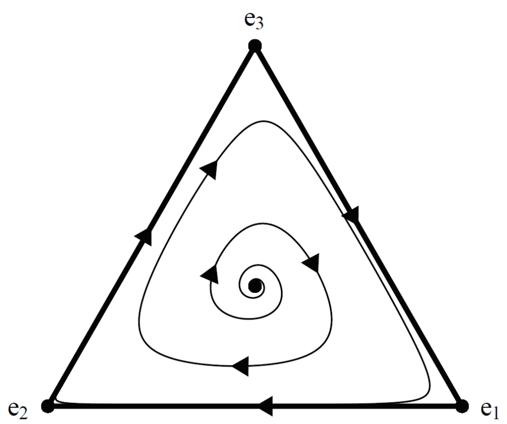

FIGURE 2. Replicator dynamics for the rock-paper-scissor game with $a_i = 1$ and $b_i = 0{,}55$ for all $i = 1, 2, 3$.

The concept of permanence means, roughly, that if all strategies are present in the population at time zero, then they will not go extinct (in the long run). Theorem 3.2 says that the time average of the solution curves of a permanent replicator dynamics equals to its unique interior stationary point.

We are now going to introduce the basic solution concept of evolutionary game theory, namely that of *evolutionary stable strategy* or, more commonly, *ESS*.

Definition 6. Consider a 2-player symmetric normal form game with strategy set Σ and payoff-function $\mathscr{J} : \Sigma \times \Sigma \to \mathbb{R}$. A mixed strategy $\mu^* \in \mathscr{M}_1(\Sigma)$ is an ESS if either

$$\mathscr{J}(\mu^*, \mu^*) > \mathscr{J}(\mu, \mu^*)$$

for all $\mu \in \mathscr{M}_1(\Sigma)\backslash\{\mu^*\}$ or else, if equality in the above for some $\hat{\mu}$,

$$\mathscr{J}(\mu^*, \hat{\mu}) > \mathscr{J}(\hat{\mu}, \hat{\mu}).$$

The notion of ESS was first introduced in [**MSP**] by Maynard Smith and Price as an alternative to Nash-equilibrium when trying to apply game theory to problems in evolutionary biology. As a concept the ESS is slightly weaker than a strict Nash-equilibrium, but is nevertheless an equilibrium in the usual sense. Intuitively, for a strategy to qualify as evolutionary stable it should, if played by all agents in a population, be resistent to attempts of invasion by any other strategy. In connection to the replicator dynamics an ESS is always an asymptotically stable stationary point and moreover, if it is an interior point of X it is even globally stable.

As we mentioned earlier another interesting path to follow in evolutionary games is that of stochastic evolutionary dynamics, being a "high-resolution" version of the mean dynamics. Indeed, according to Kurtz's theorem (see [**K3**]), we have that

$$\lim_{N\to\infty} \mathbb{P}\left(\sup_{t\in[0,T]} \left| X_N(t) - x_t \right| < \varepsilon \right) = 1$$

for any positive $T < \infty$ and $\varepsilon > 0$. Apart from questions related to convergence of population states the stochastic dynamical approach is also well suited to address problems of equilibrium selection in games with multiple locally stable equilibria. Even though interesting we will not bring up any of these results in this text, but readily refer to [**S1**] for a short survey and to [**S3**] for a more thorough discussion.

6. Zero-Sum Differential Games

In the aftermath of World War II it was clear that the destructive forces of modern warfare had evolved in a direction where advances in technology and cutting edge science were of uttermost importance. One of the most decisive (as well as disturbing) examples of such development was undoubtedly the construction of the atomic bomb, but other achievements like for instance RADAR systems, jet engine technology and cryptography also had a major impact for success in the battle field. As a consequence the RAND Cooperation (Research ANd Development) was founded in 1948, mainly by high ranked representatives from the U.S. military forces, as a think tank to "*further and promote scientific, educational, and charitable purposes, all for the public welfare and security of the United States of America*".

The RAND Cooperation attracted a rich variety of prominent researches in areas ranging from mathematics, physics and engineering to economics, psychology and philosophy. Among them were Rufus P. Isaacs who joined the mathematics department in 1948. In parallel with Richard E. Bellman (fellow researcher at RAND) Isaacs was working on problems related to optimal control and is by now often credited as founder of the subject *differential games* (see [**B2**]).

In layman terms differential games concerns the study of decision making and conflict scenarios in continuous time systems driven by differential equations. Early examples of such models are two player zero-sum games of so called pursuit-evasion type, where one player is to capture the other player escaping. Most famous is probably *The Homicidal Chauffeur Game* introduced by Isaacs in one of his RAND papers from 1951 (internal RAND document P-257) as:

"*Game I. The Homicidal Chauffeur. P and E move in the plane with speeds v_1 and v_2 with $v_1 > v_2$. The radius of curvature of P's path must never be less than a given fixed R. Time of capture is the payoff.*" ... "*For a prototype, think of a car whose driver is intent on running down an objecting pedestrian in a large unobstructed parking lot with the latter's only hope being delay until rescue.The capture region is the planform of the car.*"

It has later been revealed, according to [**B2**], that the Homicidal Chauffeur originated in a problem posed by the US Navy to RAND in the 1940's. In that version, however, the chauffeur was a torpedo and the pedestrian was a ship.

Isaacs collected his ideas and results on differential games in the, by now, classic text "*Differential games. A mathematical theory with applications to warfare and pursuit, control and optimization*" published in 1965 (see [**I**]).

To mathematically formalize the idea of a general 2-player zero-sum differential game requires two main ingredients; a *dynamics* and a *payoff function*. The dynamics describes how the state of the game evolves in time, most commonly as an ODE:

$$\begin{cases} X_t' = f(t, X_t, u_t, v_t), & t \in \mathscr{T} \\ X_{t_0} = x_0, \end{cases} \tag{6.1}$$

for a fixed initial position $(t_0, x_0) \in \mathscr{T} \times \mathbb{R}^N$ and subset $\mathscr{T} \subset \mathbb{R}_+ \cup \{\infty\}$ where $u : \mathscr{T} \to U$ and $v : \mathscr{T} \to V$ represent the controls of Player 1 and Player 2 respectively (both are assumed to be Lebesgue measurable). It is common to assume U and V to be compact metric spaces and the function f to be bounded and continuous in all of its variables and Lipschitz continuous in X_t.

For the payoff function a general form is $\mathscr{J}(\mathscr{T}, X(\cdot), u(\cdot), v(\cdot))$ where we adopt the convention that Player 2 is a maximizer and Player 1 is a minimizer. For the Homicidal Chauffeur game starting at $t_0 = 0$ with $X_t = (x_C(t), x_P(t))$, having x_C and x_P as the states of the pedestrian and the chauffeur respectively, we would take

$$\mathscr{J}(X(\cdot), u_C(\cdot), u_P(\cdot)) := \inf\{t \in \mathbb{R}_+ : x_C(t) = x_P(t)\}.$$

Another important class of payoff functions are the finite horizon *Bolza problems* (i.e. $\mathscr{T} = [t_0, T]$ for some $T > 0$), where

$$\mathscr{J}(t_0, x_0, u(\cdot), v(\cdot)) := \int_{t_0}^{T} g(s, X_s, u_s, v_s)ds + h(X_T). \tag{6.2}$$

Just as for the function defining the state dynamics g and h are often assumed to be bounded and continuous in all variables and Lipschitz continuous in the state variable. For the remainder of this introduction we limit ourselves to only consider the Bolza problem.

One of the most subtle points of differential games as compared to the classical theory of matrix games is the notion of strategy. Here we have introduced the control functions u and v as representatives of the possible actions that each player can make. Therefore, by introducing the space $\mathscr{U}(t_0)$ of Lebesgue measurable controls $u : [t_0, T] \times \mathbb{R}^N \to U$ for Player 1 and, analogously, $\mathscr{V}(t_0)$ for Player 2 a naive definition of strategy would simply be any choice of elements in $\mathscr{U}(t_0)$ and $\mathscr{V}(t_0)$. Such strategies are often called *open-loop strategies*. This notion is, however, unsatisfactory for two obvious reasons. Firstly it does not respect the fact that the dynamics in

(6.1) should be solvable regardless the choice of (u, v) and, secondly, it does not take into account that the strategy chosen by one player may affect the choice of the other. Given the regularity of f as above the first issue is readily fixed by requiring boundedness (in addition to measurability) of u and v. In more general situations the conditions required of the controls have to be sharpened. Any pair (u, v) such that (6.1) allows for a unique solution is called *admissible* and we denote the corresponding spaces of admissible controls by $\bar{\mathscr{U}}(t_0)$ and $\bar{\mathscr{V}}(t_0)$.

To formalize the idea that strategies (as realized by certain controls) may depend on each other requires some more work. Dependence of strategies naturally arises when information of one players actions is provided to the other. In general each player may have access to either *full-* or *partial information* (including no information), in equal or unequal amounts. If Player 1 and Player 2 receive the same type of information in the same amount the informational structure is called *symmetric*, otherwise *asymmetric*. The Homicidal Chauffeur is a typical example of a game providing full information to both players whereas certain Bolza type games like the Principal-Agent problem are asymmetric in partial information.

The rich variety of possible models makes a general approach towards strategies in differential games impossible. To proceed we aim at formalizing a situation in which players play in continuous time and observe each other continuously. In such a setting it is reasonable to think of strategies of Player 1 and Player 2 respectively as maps $\alpha : \bar{\mathscr{V}}(t_0) \to \bar{\mathscr{U}}(t_0)$ and $\beta : \bar{\mathscr{U}}(t_0) \to \bar{\mathscr{V}}(t_0)$. That is, from the point of view of Player 1; a strategy α is a plan for how to respond to any admissible control v of Player 2. Compared to the feedback strategies that allowed for no dependence this notion is often instead too ambitious. Intuitively we would like to think of a strategy of one player as a *nonanticipative* response to the actions of the other player, i.e. a response that does not rely on the future action of the opponent. The following definition was suggested by Varaiya in 1967 (see [**V**]) and is a good step in the right direction:

Definition 7 (Nonanticipative Strategy)**.** A map $\alpha : \bar{\mathscr{V}}(t_0) \to \bar{\mathscr{U}}(t_0)$ is a nonanticipative strategy (of Player 1) if, for any time $t_1 > t_0$ and any controls $v_1, v_2 \in \mathscr{V}(t_0)$ such that $v_1 = v_2$ almost everywhere in $[t_0, t_1]$ we have that $\alpha(v_1) = \alpha(v_2)$ almost everywhere in $[t_0, t_1]$.

The above definition is indeed appealing and it has been extensively used in the literature of differential games. It does however carry a defect of practical nature with it. When thinking of strategies as response plans one would like to have a result saying that for any pair (α, β) of nonanticipative strategies there exists a unique pair $(u, v) \in \bar{\mathscr{U}}(t_0) \times \bar{\mathscr{V}}(t_0)$ such that $\alpha(v) = u$ and $\beta(u) = v$. As a consequence one would then be able to phrase the game in normal form (i.e. express the payoff function in terms of strategies rather than controls). Such a conclusion, however, is unfortunately not true for nonanticipative strategies‡. To get around this problem one can introduce the more restrictive notion of *delay strategies* (see for instance [**C1**]):

Definition 8 (Delay Strategy). A delay strategy (of Player 1) is a nonanticipating strategy $\alpha : \bar{\mathscr{V}}(t_0) \to \bar{\mathscr{U}}(t_0)$ for which there is a delay $\tau > 0$ such that, for any two controls $v_1, v_2 \in \bar{\mathscr{V}}(t_0)$ and any $t_1 > t_0$, if $v_1 = v_2$ a.e. on $[t_0, t_1]$, then $\alpha(v_1) = \alpha(v_2)$ a.e. on $[t_0, t_1 + \tau]$.

With the rigorous notion of delay strategy at hand we now continue the study of the Bolza problem. For future reference we denote the set of delay strategies for Player 1 and Player 2 by $\mathscr{A}_d(t_0)$ and $\mathscr{B}_d(t_0)$.

In analogy with the matrix games in Section 3 it is natural to introduce the concept of value also for two player zero-sum differential games. For the Bolza problem with payoff function $\mathscr{J}$ as in (6.2) we introduce the *upper value function* as

$$\mathbf{V}^+(t_0, x_0) := \inf_{\alpha \in \mathscr{A}_d(t_0)} \sup_{\beta \in \mathscr{B}_d(t_0)} \mathscr{J}(t_0, x_0, \alpha, \beta),$$

and the *lower value function* as

$$\mathbf{V}^-(t_0, x_0) := \sup_{\beta \in \mathscr{B}_d(t_0)} \inf_{\alpha \in \mathscr{A}_d(t_0)} \mathscr{J}(t_0, x_0, \alpha, \beta).$$

By construction $\mathbf{V}^-(t_0, x_0) \leq \mathbf{V}^+(t_0, x_0)$ for all $(t_0, x_0) \in [t_0, T] \times \mathbb{R}^N$ so the task is to prove equality and characterize the value function $\mathbf{V} = \mathbf{V}^+ = \mathbf{V}^-$. A key result for doing this is stated in the following theorem (see [**C1**]):

Theorem 6.1 (Dynamic Programming). *Let $(t_0, x_0) \in [0, T) \times \mathbb{R}^N$ and $\tau \in (0, T - t_0)$. Then*

$$\mathbf{V}^+(t_0, x_0) = \inf_{\alpha \in \mathscr{A}_d(t_0)} \sup_{v \in \bar{\mathscr{V}}(t_0)} \int_{t_0}^{t_0+\tau} g(s, X_s, \alpha(v)_s, v_s) ds + \mathbf{V}^+(t_0 + \tau, X_{t_0+\tau}).$$

The analogous result for $\mathbf{V}^-(t_0, x_0)$ also holds.

‡Take $U = V = [-1, 1]$ and consider the pair of nonanticipating strategies defined by $\alpha(v)(t) = v(t)$ and $\beta(u)(t) = u(t)$ for a.e. $t \geq t_0, \forall (u, v) \in \bar{\mathscr{U}}(t_0) \times \bar{\mathscr{V}}(t_0)$. Then all controls $(u, v) \in \mathrm{Diag}(\bar{\mathscr{U}}(t_0) \times \bar{\mathscr{V}}(t_0))$ satisfy that $\alpha(v) = u$ and $\beta(u) = v$.

The conclusion in Theorem 6.1 is the well known principle of dynamic programming. It has two important consequences. First of all it implies boundedness and Lipschitz continuity of $\mathbf{V}^+$ and $\mathbf{V}^-$ in all variables. This in turn implies almost everywhere differentiability by Rademacher's theorem. Secondly it suggests a possible way characterizing $\mathbf{V}$ as the solution to a certain PDE. In particular, by the dynamic programming principle it is an easy exercise prove (at least heuristically) that $\mathbf{V}^+$ should satisfy the equation

$$\left\{ \begin{array}{l} \partial_t \phi(t,x) + \mathscr{H}^+(t,x,D\phi(t,x)) = 0, \text{ in } (0,T)\times\mathbb{R}^N \\ \phi(T,x) = g(x), \end{array} \right. \tag{6.3}$$

where $\mathscr{H}^+(t,x,p) = \inf_{u\in U}\sup_{v\in V}\{\langle p, f(t,x,u,v)\rangle + l(t,x,u,v)\}$. The analogous equation for $\mathbf{V}^-$ reads

$$\left\{ \begin{array}{l} \partial_t \phi(t,x) + \mathscr{H}^-(t,x,D\phi(t,x)) = 0, \text{ in } (0,T)\times\mathbb{R}^N \\ \phi(T,x) = g(x), \end{array} \right. \tag{6.4}$$

where $\mathscr{H}^-(t,x,p) = \sup_{v\in V}\inf_{u\in U}\{\langle p, f(t,x,u,v)\rangle + l(t,x,u,v)\}$. Thus, assuming *the Isaacs' condition*[§], i.e. $\mathscr{H}^+ = \mathscr{H}^- =: \mathscr{H}$, there is hope to have existence of a value function since $\mathbf{V}^+$ and $\mathbf{V}^-$ then would solve the same problem. In that case, given that (6.3) and (6.4) admit unique solutions, one should be able to find $\mathbf{V} = \mathbf{V}^+ = \mathbf{V}^-$ as the solution to the PDE

$$\left\{ \begin{array}{l} \partial_t \mathbf{V}(t,x) + \mathscr{H}(t,x,D\mathbf{V}(t,x)) = 0, \text{ in } (0,T)\times\mathbb{R}^N \\ \mathbf{V}(T,x) = g(x), \end{array} \right. \tag{6.5}$$

that is, the celebrated *Hamilton-Jacobi-Isaacs equation*. Despite the elegance of this result it is unfortunately not as clear characterization of the value as it seems. Indeed one can check that $\mathbf{V}^+$ solves (6.3) (and similar for $\mathbf{V}^-$) at all points of differentiability. It is however well known that equations of Hamilton-Jacobi type, in general, do not admit uniqueness in the classical sense, as the following example from [**C1**] shows:

Example 3. Consider the Hamilton-Jacobi equation

$$\left\{ \begin{array}{l} \partial_t \mathbf{V}(t,x) + |D\mathbf{V}(t,x)| = 0, \text{ for } (t,x)\in[0,T]\times\mathbb{R} \\ \mathbf{V}(T,x) = 0. \end{array} \right.$$

Then all functions of the form

$$\mathbf{V}(t,x) = \left\{ \begin{array}{l} 0, \text{ if } |x| \geq T-t, \\ k(T-t) - k|x|, \text{ if } |x| < T-t, \end{array} \right.$$

[§]It is a common feature in mathematics to use the definite article in combination with a genitive like "the Isaacs' condition" rather than the more grammatical form "the Isaacs condition". In [**F3**] this is described as an example of a "cast-iron idiom".

for any $k \in \mathbb{R}_+$, are Lipschitz continuous and satisfy (3) at any point of differentiability.

To overcome this difficulty Crandall and Lions in 1983 introduced the idea of *viscosity solutions* (see [**CL**]). Even though the importance of this concept is substantial, it is nonetheless a theory of its own and reaches beyond the scope of this presentation. We conclude this section with a result on existence and characterisation of a value function in the viscosity solution sense (see [**C1**]).

Theorem 6.2. *Under the common regularity conditions on f, l and g stated in this section and the Isaacs' condition*

$$\mathscr{H}^+(t,x,p) = \mathscr{H}^-(t,x,p), \text{ for all } (t,x,p) \in [0,T] \times \mathbb{R}^N \times \mathbb{R}^N$$

the game has a value

$$\mathbf{V}(t,x) = \mathbf{V}^+(t,x) = \mathbf{V}^-(t,x), \text{ for all } (t,x) \in [0,T] \times \mathbb{R}^N$$

where $\mathbf{V}$ *is the unique viscosity solution of the Hamilton-Jacobi-Isaacs equation (6.5).*

Remark 1. The Isaacs' condition may seem like an apparent but rather strong assumption to make in a general situation. Indeed, the problem of existence and characterization of a value in 2-player zero-sum differential games of Bolza type without the Isaacs' condition has recently been studied by Buckdahn, Li and Quincampoix in [**BLQ**]. By introducing a concept of mixed nonanticipative strategies, much like in the classical theory of matrix games, they are able to prove the existence of a value function and characterize it as the unique viscosity solution of a corresponding Hamilton-Jacobi-Isaacs equation. However, one should keep in mind that allowing for such strategies is a substantial relaxation of the original problem.

7. Stochastic Optimal Control Theory

In light of the presentation in Section 6, a natural generalization of the Bolza type differential games would arise by introducing random noise. More precisely, if W_t is a standard N-dimensional Brownian motion (on a fixed underlying probability space) we can consider the 2-player zero-sum *stochastic differential game*, defined by the dynamics

$$\begin{cases} dX_t = f(t,X_t,u_t,v_t)dt + \sigma(t,X_t,u_t,v_t)dW_t, & t \in [t_0,T], \\ X_{t_0} = x_0, \end{cases}$$

and the payoff function

$$\mathscr{J}(t_0, x_0, u(\cdot), v(\cdot)) := \mathbb{E}\left[\int_{t_0}^{T} g(s, X_s, u_s, v_s)ds + h(X_T)\right].$$

This problem was first studied rigorously by Fleming and Souganidis in [**FS2**] and turns out to be substantially more intricate than the deterministic case. One obstacle is the dynamic programming principle, which does not follow as easily (due to issues of measurability). To prove it one has to introduce new concepts of strategies like *r-strategy* and *π-admissible strategy* and derive the result by an indirect approach. Having the dynamic programming principle at hand the existence of a unique value function can be established, characterized as the unique viscosity solution of the associated *Bellman-Isaacs partial differential equation* (which in contrast to the Hamilton-Jacobi-Isaacs equation also contains second order spatial derivatives). As for the deterministic case, this result also relies on the Isaacs' condition. Since the publication of [**FS2**] in 1989 stochastic differential games have been studied by several authors, for instance Buckdahn, Cardaliaguet and Rainer in [**BCR**] (for the nonzero sum case), Bayraktar and Poor in [**BP**] (in an n-player non-Markovian setting with fractional Brownian noise) and others.

One of the most important subclasses of stochastic differential games is undoubtedly the "1-player case", more commonly known as *stochastic optimal control theory*, which we will consider next. For this discussion, think of the single player as a controller guiding a process X_t described by

$$\begin{cases} dX_t = f(t, X_t, u_t)dt + \sigma(t, X_t, u_t)dW_t, & t \in [0, T], \\ X(0) = x_0, \end{cases} \tag{7.1}$$

by means of choosing a proper process u_t (control) so that the *cost functional*

$$\mathscr{J}(u(\cdot)) := \mathbb{E}\left[\int_{0}^{T} g(s, X_s, u_s)ds + h(X_T)\right], \tag{7.2}$$

is minimized (here we have suppressed the dependence of x_0 in (7.2) and chosen $t_0 = 0$ for convenience). The objective is to characterize and, if possible, find optimal choices of u_t.

Even though the roots of (deterministic) optimal control theory are to be found in calculus of variations, thus dating back to the seventeenth century and the studies of Johann Bernoulli, the igniting spark of the subject is often credited the work of Lev S. Pontryagin in the 50's and his group at the USSR Steklov Institute. Their studies culminated in the important (and mathematically rigourous) characterization of optimal controls, which by now goes under the name of *Pontryagin's maximum principle*. It has, however, been pointed out by for instance Pesch and Plail [**PP**] that earlier achievements by Carathéodory and Graves in the 1930's should be acknowledged as precursors of this result. Also Magnus R. Hestenes at RAND gave a formulation of the maximum principle for Bolza type problems, slightly before Pontryagin, but his result is much weaker.

Before we go on presenting the most fundamental achievements in stochastic optimal control we need to specify the problem formulation. Especially, what type of processes u are we interested in? Assuming an underlying, complete and filtered probability space $(\Omega, \mathscr{F}, \{\mathscr{F}_t\}, \mathbb{P})$ a natural answer, which is indeed very general, would be to search among the $\{\mathscr{F}_t\}$-adapted processes taking values in some predefined subset $U \subset \mathbb{R}$ (possibly unbounded). For future reference we denote the set of such processes by $\mathscr{U}[0,T]$. For applications, however, $\{\mathscr{F}_t\}$-adaptivity is sometimes a too general assumption since, for instance, the random noise may not be observable. In such cases it is more suitable to consider processes of the form $u_t = u(t, x_{[0,t]})$ in the set of *natural controls* $\mathscr{U}_N[0,T]$, or of the form $u_t = u(t, x_t)$ in the set of *feedback* (or Markov) *controls* $\mathscr{U}_M[0,T]$. By construction $\mathscr{U}_M[0,T] \subset \mathscr{U}_N[0,T] \subset \mathscr{U}[0,T]$ so we clearly have

$$\inf_{u \in \mathscr{U}_M[0,T]} \mathscr{J}(u(\cdot)) \geq \inf_{u \in \mathscr{U}_N[0,T]} \mathscr{J}(u(\cdot)) \geq \inf_{u \in \mathscr{U}[0,T]} \mathscr{J}(u(\cdot)). \tag{7.3}$$

Regardless what type of control to consider, we must assume that for any u in the set of feasible controls, the SDE in (7.1) admits a unique solution. Even though existence theory of SDEs is well understood (see for instance [**Ø**]) this assumption is rather delicate and not at all innocent. For controls of feedback type the conditions

$$\begin{aligned}
&|f(t,x,u(t,x)) - f(t,y,u(t,y))| \leq K|x-y|, \\
&|\sigma(t,x,u(t,x)) - \sigma(t,y,u(t,y))| \leq K|x-y|, \\
&|f(t,x,u(t,x))| + |\sigma(t,x,u(t,x))| \leq K(1+|x|),
\end{aligned}$$

for some positive constant K, are sufficient for the existence of a unique $\{\mathscr{F}_t\}$-adapted solution of (7.1) with $\mathbb{P}$-a.s. continuous trajectories. We start in this setup, i.e. minimizing (7.2) over $\mathscr{U}_M[0,T]$.

As in Section 6 it is natural to introduce a value function, but here we do it in a slightly different manner:

$$V(t,x) := \inf_{u\in\mathscr{U}_M[0,T]} \mathbb{E}\left[\int_t^T g(s,X_s,u_s)ds + h(X_T)\,\middle|\, X_t = x\right]. \tag{7.4}$$

In analogy with Theorem 6.1 $V(\cdot,\cdot)$ satisfies a dynamic programming principle, leading to the important *Hamilton-Jacobi-Bellman equation* (often abbreviated the HJB-equation):

Theorem 7.1 (The HJB-equation)**.** *Assume that $V \in \mathscr{C}^{1,2}$ (i.e. continuously differentiable in t and two times continuously differentiable in x) and that $\bar{u} \in \mathscr{U}_M[0,T]$ is an optimal control. Then V satisfies the HJB-equation*

$$\begin{cases} \partial_t V(t,x) + \inf_{u\in U}\{g(t,x,u) + L^u V(t,x)\} = 0, \text{ in } (0,T)\times\mathbb{R}^N, \\ V(T,x) = h(x), \end{cases} \tag{7.5}$$

where L^u denotes the infinitesimal generator of (7.1) given u, i.e.

$$L^u := \sum_i f_i(t,x,u)\partial_i + \frac{1}{2}\sum_{i,j}(\sigma(t,x,u)\sigma(t,x,u)^\top)_{i,j}\partial^2_{ij}.$$

Furthermore, for each $(t,x) \in (0,T)\times\mathbb{R}^N$ it holds that

$$\bar{u}(t,x) \in \operatorname*{arginf}_{u\in U}\{g(t,x,u) + L^u V(t,x)\}.$$

One should note that Theorem 7.1 only serves as a characterization of optimality since the existence of $\bar{u}$ is part of the assumptions. Thankfully, the setting of the problem allows for a *Verification Theorem* (i.e. sufficient conditions for optimality) where the HJB-equation also acts as the protagonist:

Theorem 7.2 (Verification Theorem)**.** *Suppose that $\phi : (0,T]\times\mathbb{R}^N \to \mathbb{R}$ is a function such that*

$$\int_0^T \mathbb{E}\left[\nabla_x\phi(s,X_s)\sigma(s,X_s,u_s)\right]^2 ds < \infty,$$

for any choice of feedback control $u \in \mathscr{U}_M[0,T]$, and that

$$\begin{cases} \partial_t\phi(t,x) + \inf_{u\in U}\{g(t,x,u) + L^u\phi(t,x)\} = 0, \text{ in } (0,T)\times\mathbb{R}^N, \\ \phi(T,x) = h(x). \end{cases}$$

Also suppose that $v \in \mathscr{U}_M[0,T]$ is such that

$$v(t,x) = \operatorname*{arginf}_{u\in U}\{g(t,x,u) + L^u\phi(t,x)\},$$

for any fixed pair (t,x). Then the value function in (7.4) equals to ϕ and there exists an optimal control $\bar{u}$ of feedback type given by $\bar{u}(t,x)=v(t,x)$.

Despite the beauty of the HJB-equation and its consequences to optimal control, the results presented thus far carry some conceptual drawbacks. The restriction to consider controls of feedback type is often practical, but it is also rather heavy. It is therefore natural to ask oneself; is it possible to have equalities in (7.3) and, if so, under what conditions? The answer to this question is yes, but the results put further restrictions on the structure of the problem. We refer to [**K2**], [**FS1**] and [**Ø**] for a more thorough discussion.

Another inconvenience with the HJB-equation is, as we have already encountered in Section 6, the poor theory of existence and uniqueness for classical solutions to (7.5). In fact, the assumption of having $V \in \mathscr{C}^{1,2}$ is only valid in a rather narrow class of problems. In order to get a more comprehensive approach one inevitably has to dig into the technical theory of viscosity solutions. The next method that we are going to discuss, Pontryagin's maximum principle, neatly resolves both of these issues.

Again, consider the optimal control problem defined by (7.1) and (7.2), now with respect to the space of adapted controls $\mathscr{U}[0,T]$. To get the theory run smoothly it is common to require f, σ, g and h to be $\mathscr{C}^2$ with respect to x. Moreover, these functions and all their derivatives up to second order with respect to x should be continuous in (x,u), and bounded. For a more refined (and technical) discussion of regularity we refer to [**YZ**]. The goal of the maximum principle is to derive a set of necessary conditions for characterizing optimality, not relying on the principle of dynamic programming. The idea is to start from an optimal control $\bar{u}$, then construct an adjacent control $u^\varepsilon \in \mathscr{U}[0,T]$ from it and perform a perturbation analysis of $\mathscr{J}(u^\varepsilon)-\mathscr{J}(\bar{u})$ directly by means of a Taylor expansion. One of the key achievements in this enterprise is the construction of u^ε by means of *spike variation* (or *needle variation*). That is, for some $\varepsilon>0$, pick a subset $E_\varepsilon \subset [0,T]$ such that $|E_\varepsilon|=\varepsilon$ (Lebesgue measure) and consider

$$u^\varepsilon(t) := \begin{cases} u(t), t \in E_\varepsilon, \\ \bar{u}(t), t \in E_\varepsilon^c, \end{cases}$$

for an arbitrarily chosen control $u \in \mathscr{U}[0,T]$. The goal is to prove a relation of the form

$$\mathscr{J}(u^\varepsilon)-\mathscr{J}(\bar{u}) = \mathbb{E}\left[\int_0^T \delta\mathscr{H}(s)\cdot\chi_{E_\varepsilon}(s)ds\right]+o(\varepsilon),$$

where $\delta\mathscr{H}(t) := \mathscr{H}(t,\bar{X}_t,u(t),\cdot) - \mathscr{H}(t,\bar{X}_t,\bar{u}(t),\cdot)$ for a given function $\mathscr{H}$, and establish a characterization of $\bar{u}$ in the limit as $\varepsilon \to 0$. The following result was proven by Peng in 1990 (see [**P**]):

Theorem 7.3 (Peng's Stochastic Maximum Principle)**.** *For simplicity, consider the case $N = 1$ and let the previously stated regularity assumptions on f, σ, g and h hold. If $(\bar{x}(\cdot),\bar{u}(\cdot))$ is a solution of the optimal control problem (7.1), (7.2), then we have*

$$(p(\cdot),q(\cdot)),(P(\cdot),Q(\cdot)) \in L^2_{\mathscr{F}}(0,T;\mathbb{R}) \times L^2_{\mathscr{F}}(0,T;\mathbb{R}),$$

which are respectively solutions to the backward stochastic differential equations (BSDEs)

$$\left\{\begin{array}{l} dp(t) = -H_x(t)dt + q(t)dW(t), \\ p(T) = -h_x(\bar{x}(T)), \end{array}\right. \tag{7.6}$$

and

$$\left\{\begin{array}{l} dP(t) = -\left\{2f_x(t)P(t) + \sigma_x^2(t)P(t) + 2\sigma_x(t)Q(t) + H_{xx}(t)\right\}dt \\ \qquad\qquad +Q(t)dW(t), \\ P(T) = -h_{xx}(\bar{x}(T)), \end{array}\right. \tag{7.7}$$

where $H(t) := f(t)p(t) + \sigma(t)q(t) - g(t)$[¶]*, such that*

$$\begin{array}{l} H(t,\bar{x}(t),u,p(t),q(t)) - H(t,\bar{x}(t),\bar{u}(t),p(t),q(t)) \\ +\frac{1}{2}P(t)(\sigma(t,\bar{x}(t),u) - \sigma(t,\bar{x}(t),\bar{u}(t)))^2 \le 0, \end{array} \tag{7.8}$$

for all $u \in U$, a.e. $t \in [0,T]$, $\mathbb{P}$-a.s.

Remark 2. If the diffusion coefficient σ does not depend on the control variable, then (7.8) simplifies to the characterization

$$\bar{u}(t) \in \operatorname*{argmax}_{u\in U}\ H(t,\bar{x}(t),u,p(t),q(t)),$$

making the solution to (7.7) superfluous.

Remark 3. The BSDEs (7.6) and (7.7) are commonly referred to as the first- and second order *adjoint equations*. Backward stochastic differential equations were first introduced by Bismut in [**B1**] and require, in contrast to regular SDEs, solutions given as pairs of processes. The reason for this being to ensure $\{\mathscr{F}_t\}$-adaptivity, since a simple time reversal would not take the informational structure into account.

[¶]Here, for $\varphi \in \{f,\sigma,g,h\}$, we use the simplified notational convention that $\varphi(t) := \varphi(t,\bar{x}(t),\bar{u}(t))$.

In spirit of the Verification Theorem 7.2 for the HJB-equation one can also turn the maximum principle into a sufficient condition of optimality by imposing convexity assumptions on the functions h and $H(t,\cdot,\cdot,p(t),q(t))$ and on the space $U \subset \mathbb{R}$.

A natural question to ask at this point is; in what way do the HJB-equation and the maximum principle relate to each other? Indeed, the appearance of Theorem 7.5 as compared to Theorem 7.3 is seemingly different but, nevertheless, they both constitute necessary conditions for optimality of the same underlying control problem. Restricting ourselves to the deterministic case, the answer to this question is (partly) explained by a well-known construction for solving first order PDEs by means of ODEs, known as *method of characteristics*. Without digging into details (thus under suitable conditions of regularity) this result states that if there exists a control $\bar{u}(\cdot)$ such that for any $(s,y) \in [0,T] \times \mathbb{R}^N$ there exists a pair of solutions $(x^{s,y}(\cdot), p^{s,y}(\cdot))$ to the system

$$\text{(7.9)} \qquad \begin{cases} \dot{x}^{s,y}(t) = \dfrac{\partial H}{\partial p}(t), \ t \in [s,T] \\ \dot{p}^{s,y}(t) = -\dfrac{\partial H}{\partial x}(t), \ t \in [s,T] \\ x^{s,y}(s) = y, \\ p^{s,y}(s) = -h_x(x^{s,y}(T)) \end{cases}$$

and

$$\bar{u}(t) \in \operatorname*{argmax}_{u \in U} \ H(s,y,u,p^{s,y}(s)), \quad \text{for all } (s,y) \in [0,T] \times \mathbb{R}^N,$$

then the function

$$\phi(s,y) := \int_s^T g(t, x^{s,y}(t), \bar{u}(t))dt + h(x^{s,y}(T)),$$

solves the HJB-equation. Note that the Hamiltonian system in (7.9) is nothing but a delayed deterministic version of (7.1) and (7.6) from the maximum principle.

For the stochastic optimal control problem the situation is rather different. Then, since the HJB-equation is a second order PDE, a connection to the maximum principle by the method of characteristics cannot be established. However, under suitable regularity one can find interesting relations between the value function $V(t,x)$ and the adjoint processes $p(t)$ and $q(t)$. We refer to [**YZ**] for a thorough discussion of these matters and conclude this introduction by a few referential remarks.

The initial work in stochastic optimal control theory by means of the maximum principle was done by Kushner in [**K4**], published in 1972. He considered a problem in which the diffusion coefficient in (7.1) did not depend on the control variable and where the optimization was subject to a set of state constraints. This study was later generalized by Bismut [**B1**], which in turn was methodologically simplified, from stochastic convex analysis to perturbation theory, by Bensoussan in [**B2**]. Since the important contributions by Peng [**P**] the subject of stochastic maximum principles has continued to attract scientific interest, both within pure mathematical research and various fields of application. Recent contributions by Andersson and Djehiche in [**AD**] and Buckdahn, Djehiche and Li in [**BDL**] are a good examples of such a development, where Theorem 7.3 is generalized to certain problems of mean-field type. Also methods of dynamic programming and HJB-equations has ever since the studies of Bellman in the 50's (see e.g. [**B1**]) continued to break new ground, not least within mean field games and contributions by Lasry and Lions [**LL1**], [**LL1**] and Caines, Huang and Malhamé [**HMC**]. As for the fields of stochastic optimal control and game theory, the interplay between them and their relevance in applications, many challenging and important problems remain to be solved, paving for new exciting discoveries to be made.

8. Summary of papers

We end this introduction by a brief description of each paper appended to the thesis.

8.1. Paper I. The first paper of this thesis treats asymptotic properties of two different N-player models of the classical game *The War of Attrition*, introduced by Haigh and Cannings in [**HC**], as the number of players grows to infinity. By analysing the models in the limit regime we gain insights of the large scale characteristics of the games that are used to establish new results for one of the finite models.

The War of Attrition was first introduced by John Maynard Smith in 1974 in the well known paper *Theory of games and the evolution of animal conflicts* (see [**MS**]). The game considers two identical players competing for one prize V of positive value by observing each other and waiting for the opponent to quit. The first player to do so looses the competition and leaves the prize to the remaining opponent. In the classical setup there is a cost of waiting which is linear in time, i.e. by waiting t units of time the player gets the possibility of winning the prize $V > 0$, but he will also be obligated to pay t units in time cost. If Player X and Player Y choose waiting times τ_x and τ_y respectively, the payoff function of Player X can be written:

$$\mathscr{J}_X(\tau_x, \tau_y) := \begin{cases} V - \tau_y, & \text{if } \tau_x > \tau_y \\ V/2 - \tau_x, & \text{if } \tau_x = \tau_y \\ -\tau_x, & \text{if } \tau_x < \tau_y. \end{cases}$$

Note that the winning player pays the same time cost as that of the loosing player since he can observe his opponent leaving the game.

Equilibrium analysis of the 2-player game was done by Bishop and Cannings in 1976 who proved that the War of Attrition admits a unique ESS in mixed strategies, given by an exponential distribution of mean V (see [**BC**]).

In 1989 Haigh and Cannings constructed two canonical N-player generalisations of The War of Attrition; *the dynamic model* and *the static model*. The dynamic model is an N-player repetitive game having a sequence of prizes $\{V_k\}_{k=1}^N \subset \mathbb{R}_+$ at stake and is played in $N-1$ rounds. The first round begins by letting all players choose individual waiting times (independently of each other). The player having the least waiting time wins the prize V_1, pays his time cost (still linear in time) and leaves the game. The players remaining pay the same time cost as the player leaving and enter the second round which proceeds just as the first, having the prize V_2 at stake. The final $(N-1)$'th round thus becomes a normal 2-player War of Attrition. For increasing prize sequences, i.e. $V_1 < V_2 < ... < V_N$, it is proven in [**HC**] that the dynamic model, just as the 2-player War of Attrition and in each round k, admits a unique mixed ESS given by an exponential distribution, but with mean $(N-k)(V_{k+1} - V_k)$. Also the general case of arbitrary sequences is analysed, still having the existence of a unique ESS as a result (however not as explicit as in the case of increasing sequences).

The static model differs from the dynamic model in being a one-shot game rather than a repetitive, that is, the game finishes in one turn. Just as in the dynamic model the static model starts by letting each participating player pick a waiting time (independently). The results are then presented and prizes are handed out in the natural order, i.e. the player having the least waiting time receives V_1, the player having the second least receives V_2 and so forth. All players pay their individual time cost except for the "last" player who pays the time cost of the second last player (so that, for $N = 2$, we get back to the original 2-player War of Attrition).

The equilibrium analysis of the static model is a bit more intricate than in the dynamic model since all players, in a sense, are competing for all prizes in $\{V_k\}_{k=1}^N$ at once. In [**HC**] it is proven that the static model admits a unique ESS for prize sequences such that $V_{k+1} - V_k = c > 0$. For more general sequences however, the question of existence and uniqueness is unclear. For instance, the 3-player game generated by the prize sequence $\{1, 4, 6\}$ admits a unique mixed Nash-equilibrium which is not an ESS. There are even games in the static model that lack Nash-equilibria, like for example the game generated by $\{1, 2, 1\}$. The goal of this paper is to study asymptotic behaviors in the dynamic- and the static model of the N-player War of Attrition as $N \to \infty$.

In Section 2 of the paper we present a heuristic approach for analysing the limiting behaviour of the dynamic model as the number of players tend to infinity. To maintain regularity in the limit we introduce the concept of *prize function* $V(x)$, defined on the compact unit interval, to replace prize sequences by making the assumption that $V_k = V(k/N)$, $1 \le k \le N$. We also assume $V(x)$ to be an increasing function in $\mathscr{C}^1[0,1]$. In the limit we find that the fraction of players q that has left the game at time t (after game start at $t = 0$) is given by the equation $V(q(t)) = t - V(0)$ and, in particular, if $V(0) = 0$ then $q(t) = V^{-1}(t)$.

The results from Section 2 are rigourously studied in Section 3 where we proceed with the asymptotic analysis of the dynamic model. Considering N players we introduce the continuous time Markov chain $X(t)$ (suppressing the index N) counting the fraction of players that have left the game at time t. Given a prize function $V(x)$ such that $V(0) = 0$ it is, by the results from Section 2, natural to believe that "$X \to V^{-1}$" in some sense. Indeed, in Corollary 3.3 we prove that $\lim_{N\to\infty} \mathbb{E}[X(t)] = V^{-1}(t)$ and $\lim_{N\to\infty} \mathrm{Var}(X(t)) = 0$ on the time interval $t \in [0, V(1))$. In addition to the result of convergence in mean of $X(t)$ we also manage to prove in Theorem 3.6 that (in a sense) the limit dynamic model behaves like a static model with players using the mixed strategy $\dot{q}(t) = d/dt(V^{-1})(t)$. We therefore have a good reason to proceed by analysing the asymptotic properties of the static model.

Section 4 is devoted to convergence properties in the static model of the War of Attrition. In [**HC**] one can find a necessary condition for a smooth probability density to be a static model ESS, stated as a nonlinear autonomous ODE-problem of its CDF. For a general prize sequence $\{V_k\}_{k=1}^N$ this

ODE might be singular, but under the assumption of monotonicity (increasing) the right hand side is always well defined. By well known properties of asymptotic stability in autonomous equations we give an argument (valid for any increasing prize sequence) for the existence and uniqueness of a solution and for why the solution is the CDF of a probability density. Furthermore, by introducing a prize function $V(x)$ just as we did in Section 2 and Section 3, we are able to prove Proposition 4.2 saying that the solution converges uniformly to $q(t) = V^{-1}(t)$ on $t \in [0, V(1))$ as $N \to \infty$. Hence the dynamic- and the static model "coincide" in the limit of infinitely many players.

Section 4 explained to us that the only candidate ESS in the N-player limit of the static model is given by the density function $\dot{q}(t) = d/dt(V^{-1}(t))$. In Section 5 we analyse wether $\dot{q}(t)$ is an ESS or not and establish new results for the static model in finitely many players. By introducing theory of normal form games with a continuum of players (according to [**B**]) we define the notion of ESS (Appendix C). Assuming the prize function to be in $\mathscr{C}^2[0,1]$ rather than in $\mathscr{C}^1[0,1]$ (still increasing and normalised so that $V(0) = 0$) it is an easy task to prove by direct calculation that $\dot{q}(t)$ is an ESS in the continuum model if and only if $V(x)$ is strictly convex. Moreover, if $V(x)$ is concave the $\dot{q}$-strategy is not an ESS. The importance of convexity/concavity of $\{V_k\}_{k=1}^N$ is not obvious in the N-player static model, but it is reasonable to believe that the conclusions made in the continuum limit also should hold in the finite case for large N. Surprisingly enough we can prove as a corollary to Theorem 5.3 that the static model does admit a unique ESS, not only for N large enough, but for for all $N \geq 2$ if the prize sequence (not necessarily connected to a prize function) is convex. The concave case turns out to be a bit more difficult, but in Theorem 5.5 we manage to prove that if $V(x) := x^\alpha$, for any $0 < \alpha < 1$, (hence V is concave) and $V_k = V(k/N)$, then, for any N large enough, the static model lacks an ESS.

8.2. Paper II. This paper presents an approach for characterizing optimal contracts in a continuous time *Principal-Agent problem* (from here on PA). Using the powerful machinery of Pontryagin's maximum principle we are able to consider a very general class of models, for instance non-Markovian settings where the dynamic programming principle fails to be fulfilled.

The main source of inspiration for much of recent years progress in continuous time versions of PA is Holmström and Milgrom [**HM**]. They rigourously study a discrete time model in T periods involving two participants, one agent and one principal. For simplicity we refer to them as A and P respectively. The goal of P is to hire A for managing a project delivering a certain observable (noisy) output x_i after each period $i \in \{1, 2, \dots, T\}$. To compensate A for his/her efforts P suggests an incentive scheme (salary) paying $s(X_T)$ at the end of the last period. Here $X_T = (x_1, \dots, x_T)$, i.e. the history of output. An important feature is the structure of s, which is supposed to depend on output only, and not on A's effort which is considered hidden information. Both A and P are supposed value wealth/cost according to exponential utility functions with fixed (but different) absolute risk aversions. Furthermore, A will only accept s and start working for P if it meets a certain *participation constraint*, in short a *contract*, ensuring A the possibility to earn a fixed minimal amount (utility wise) from the project. The problem of P is to tailor an incentive scheme and suggest a work load (that is; how much effort to put into the project) which mutually optimizes the utility functions of A and P, and respects the participation constraint. A celebrated result in [**HM**] is a theorem characterizing optimality of s by linearity (see Theorem 5 p. 314).

Inspired mainly by Williams [**W**] and Cvitanić and Zhang [**CZ**], both studying continuous time versions of [**HM**] using a maximum principle approach, Paper II employs a sequential optimization scheme, divided into *the Agent's problem* and *the Principal's problem*, to characterize optimal contracts. The Agent's problem is about characterizing optimal choice of effort $e(\cdot)$ and the Principal's problem characterizes optimal cash-flow (or incentive scheme) $s(\cdot)$. In this context we consider two different settings of the PA problem; *the Hidden Action model* and *the Hidden Contract model*.

The common feature in both of these models, from a methodological point of view, is that the Agent's- and the Principal's problem both can be formulated as nonlinear problems of stochastic optimal control, coupled to each other by the agent's effort and the continuously running cash-flow as separate controls. The Agent's problem is standard, having a dynamics (i.e. output) defined by an SDE, whereas the Principal's problem is a state constrained control problem with an FBSDE-dynamics.

The Hidden Action model is a natural generalization of the discrete time model in [**HM**]. Here the principal suggests a cash-flow $s(\cdot)$ adapted to the filtration generated by output x. In order to get the Agent's problem mathematically tractable we consider a Girsanov transformation and, in line with

[**W**] and [**CZ**], declare the *density of output* (i.e. the Radon-Nikodym derivative of the involved probability measures) $\Gamma(\cdot)$ as the controlled dynamics rather than $x(\cdot)$. We can then characterize optimal effort of the agent by a standard stochastic maximum principle argument. The next step is to insert this characterisation into the Principal's problem, which is slightly more difficult for a couple of reasons. First of all, the dynamics to control is a FBSDE-system consisting of the density of output SDE coupled to the first order adjoint equation from the Agent's problem. Secondly, the problem is constrained due to the participation constraint. In order to handle such a control problem we consult recent results in the theory of FBSDEs and prove the state constrained version of the maximum principle. A full characterization of optimality in the Hidden Action model is then established in Section 3.

The Hidden Contract model can be seen as a relaxation of the Hidden Action model. By enlarging the informational structure of the principal, from the filtration generated by output to the filtration generated by the Brownian noise, the Agent's problem becomes a feasible control problem even without applying a Girsanov transformation. Accepting such a setup we are able to state a full characterization of optimality in the same manner as described above, that is, a sequential chain of control problems. We also consider an explicit example in a linear quadratic setting which we are able to fully solve. It turns out that the optimal contract is characterized by processes which are adapted to the filtration generated by $x(\cdot)$.

In Section 5 we have made an attempt to consider a Hidden Action model in the strong formulation, i.e. without the Girsanov transformation. To get a tractable problem we restrict ourselves to consider cash-flows of feedback type $s(t, x(t))$. In particular we study an interesting example of linear quadratic nature. It turns out that optimal effort in the Agent's problem then is characterized as the unique solution to an associated nonhomogeneous Burgers' equation. Such equations admit explicit solutions by an ingenious transformation known as the Hopf-Cole transformation. The resulting problem for the principal, however, is a non standard problem of optimal control and resembles the intractable Agent's problem in the full Hidden Action model. We have neither been able to solve this problem, nor have we been able to derive a clear characterization of optimality. The problem has therefore been left for a future study.

An Open Problem. As presented in Section 5 of Paper II (Example 5.1) the unique characterization of optimal effort in the Hidden Action model with a feedback type cash-flow is

$$\varphi(t,x,s.) = \sigma^2 \frac{\int_t^T \int_{\mathbb{R}} \frac{\xi - x}{\sigma^2(\tau - t)} G(x,\xi,\tau - t)s(\tau,\xi)d\xi d\tau}{1 + \int_t^T \int_{\mathbb{R}} G(x,\xi,\tau - t)s(\tau,\xi)d\xi d\tau},$$

where G denotes the heat kernel. The Principal's problem then reads as follows:

Problem. Minimize

$$\mathscr{J}_P(s(\cdot)) := \mathbb{E}\left[\int_0^T s(t)dt - x(T)\right],$$

over the set of feedback controls s, subject to the state constraint

$$\mathbb{E}\left[\int_0^T \frac{\varphi(t,x(t),s.)^2}{2} - s(t,x(t))dt\right] < W_0, \tag{8.1}$$

for some $W_0 < 0$, and the dynamics

$$\begin{cases} dx(t) = \varphi(t,x(t),s.)dt + \sigma dW_t, \\ x(0) = 0. \end{cases} \tag{8.2}$$

It is for the moment unclear wether the problem stated above has got a solution or not. Even though it seems difficult, or even impossible, to find an explicit solution, it would be interesting to find a good characterisation of optimality. One approach would certainly be to consider a weak formulation in which the stochastic dynamics (8.2) is exchanged for the Fokker-Planck equation describing the time evolution of the density of x. The problem then becomes a deterministic optimal control problem with respect to a PDE. Following the lines of Bensoussan [**B**] (and neglecting the state constraint (8.1)) one indeed finds a necessary condition for optimality as

$$\varphi_s(t,x,s.)(v) \geq -\frac{v(t,x)}{\partial_x p(t,x)},$$

for any v in a proper set of test functions. Here φ_s denotes the Gâteaux derivative of φ with respect to s and p is the solution to a corresponding adjoint equation, namely

$$\begin{cases} \partial_t p(t,x) = -\frac{\sigma^2}{2}\partial_x^2 p(t,x) + (s(t,x) - \partial_x \varphi(t,x,s.)), \\ p(T,x) = x. \end{cases}$$

This characterization of optimality is, however, very complex and nonintuitive. It would be desirable to find a more practical condition, for instance in the form of a PDE. The existence of such a characterization still remains an open question.

8.3. Paper III. A modeling disadvantage in the Hidden Contract setup, where the agent merely reacts to a cash-flow adapted to the noise, is the possibility of being exposed to high levels of risk, despite an acceptable participation constraint. Questions of risk can also be brought up naturally in the Hidden Action model, for instance by declaring the principal risk averse. Paper III is a follow-up paper of Paper II where we address these problems and find necessary conditions of optimal contracts by means of stochastic maximum principles. This has, to the best of our knowledge, not been studied before in the context of PA problems.

As our main theoretical reference we use [**BDL**] in which a so called general maximum principle of mean-field type is found. Just as in Paper III we can extend this result to a state constrained FBSDE-system, also of mean-field type, and formulate the Agent's problem and the Principal's problem separately. In the Hidden Action model, as compared to Paper II, we generalize the cost functional of the principal to a mean-variance utility with a given constant risk aversion parameter. To remedy the agent's exposure to risk in the Hidden Contract model we extend the participation constraint to involve a statement of an upper limit for acceptable variance. The main results of the paper are presented in Section 4 as maximum principles for optimality characterization. We also consider an explicit example in Section 5 of Hidden Contract type which we are able to solve completely.

References

[AD] D. Andersson and B. Djehiche, *A maximum principle for SDEs of mean-field type*, Appl. Math. Optim. **63** (2011), no. 3, 341–356.

[B] E. J. Balder, *A unifying approach to existence of Nash equilibria*, Internat. J. Game Theory **24** (1995), no. 1, 79–94.

[BP] E. Bayraktar and H. V. Poor, *Stochastic differential games in a non-Markovian setting*, SIAM J. Control Optim. **43** (2005), no. 5, 1737–1756.

[B1] R. Bellman, *Dynamic programming*, Princeton Landmarks in Mathematics, Princeton University Press, Princeton, NJ, 2010. Reprint of the 1957 edition; With a new introduction by Stuart Dreyfus. MR2641641 (2011b:90002)

[B2] A. Bensoussan, *Lectures on stochastic control*, Nonlinear filtering and stochastic control (Cortona, 1981), Lecture Notes in Math., vol. 972, Springer, Berlin-New York, 1982, pp. 1–62.

[BFY] A. Bensoussan, J. Frehse, and P. Yam, *Mean field games and mean field type control theory*, Springer Briefs in Mathematics, Springer, New York, 2013.

[B] A. Bensoussan, *Stochastic maximum principle for distributed parameter systems*, J. Franklin Inst. **315** (1983), no. 5-6, 387–406.

[BC] D. T. Bishop and C. Cannings, *Models of animal conflict*, Adv. Appl. Prob. **8** (1976), 616–621.

[B1] J.-M. Bismut, *An introductory approach to duality in optimal stochastic control*, SIAM Rev. **20** (1978), no. 1, 62–78.

[B2] M. H. Breitner, *The genesis of differential games in light of Isaacs' contributions*, J. Optim. Theory Appl. **124** (2005), no. 3, 523–559.

[BCR] R. Buckdahn, P. Cardaliaguet, and C. Rainer, *Nash equilibrium payoffs for nonzero-sum stochastic differential games*, SIAM J. Control Optim. **43** (2004), no. 2, 624–642 (electronic).

[BDL] R. Buckdahn, B. Djehiche, and J. Li, *A general stochastic maximum principle for SDEs of mean-field type*, Appl. Math. Optim. **64** (2011), no. 2, 197–216.

[BLQ] R. Buckdahn, J. Li, and M. Quincampoix, *Value function of differential games without Isaacs conditions. An approach with nonanticipative mixed strategies*, Internat. J. Game Theory **42** (2013), no. 4, 989–1020.

[HMC] M. Huang, R. P. Malhamé, and P. E. Caines, *Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle*, Commun. Inf. Syst. **6** (2006), no. 3, 221–251.

[C1] P. Cardaliaguet, *Introduction to differential games*, Pierre Cardaliaguet's Home Page, Université Paris-Dauphine [Online], https://www.ceremade.dauphine.fr/∼cardalia/ . Acessed April 1 2015. (2010).

[C2] J. L. Casti, *Five golden rules: great theories of 20th-century mathematics and why they matter*, Wiley & Sons, 1996.

[CL] M. G. Crandall and P.-L. Lions, *Viscosity solutions of Hamilton-Jacobi equations*, Trans. Amer. Math. Soc. **277** (1983), no. 1, 1–42.

[CZ] J. Cvitanić and J. Zhang, *Contract theory in continuous-time models*, Springer Finance, Springer, Heidelberg, 2013.

[FS1] W. H. Fleming and H. M. Soner, *Controlled Markov processes and viscosity solutions*, 2nd ed., Stochastic Modelling and Applied Probability, vol. 25, Springer, New York, 2006.

[FS2] W. H. Fleming and P. E. Souganidis, *On the existence of value functions of two-player, zero-sum stochastic differential games*, Indiana Univ. Math. J. **38** (1989), no. 2, 293–314.

[F1] F. Forgó, *Contribution of Hungarian mathematicians to game theory*, A panorama of Hungarian mathematics in the twentieth century. I, Bolyai Soc. Math. Stud., vol. 14, Springer, Berlin, 2006, pp. 537–548.

[F2] F. Forgó, *John von Neumann's Contribution to Modern Game Theory*, Acta Oeconomica **54** (2004), 73–84.

[F3] H. W. Fowler, *A Dictionary of Modern English Usage*, Oxford Univ. Press, New York, 1965.

[F4] M. Fréchet, *Emile Borel, initiator of the theory of psychological games and its application*, Econometrica **21** (1953), 95–96.

[HC] J. Haigh and C. Cannings, *The n-person war of attrition*, Acta Appl. Math. **14** (1989), no. 1-2, 59–74. Evolution and control in biological systems (Laxenburg, 1987).

[HS] J. Hofbauer and K. Sigmund, *Evolutionary game dynamics*, Bull. Amer. Math. Soc. (N.S.) **40** (2003), no. 4, 479–519.

[HM] B. Holmström and P. Milgrom, *Aggregation and linearity in the provision of intertemporal incentives*, Econometrica **55** (1987), no. 2, 303–328.

[I] R. Isaacs, *Differential games. A mathematical theory with applications to warfare and pursuit, control and optimization*, John Wiley & Sons, Inc., New York-London-Sydney, 1965.

[K1] T. H. Kjeldsen, *John von Neumann's conception of the minimax theorem: a journey through different mathematical contexts*, Arch. Hist. Exact Sci. **56** (2001), no. 1, 39–68.

[K2] N. V. Krylov, *Controlled diffusion processes*, Stochastic Modelling and Applied Probability, vol. 14, Springer-Verlag, Berlin, 2009. Translated from the 1977 Russian original by A. B. Aries; Reprint of the 1980 edition.

[K3] T. G. Kurtz, *Solutions of ordinary differential equations as limits of pure jump Markov processes*, J. Appl. Probability **7** (1970), 49–58.

[K4] H. J. Kushner, *Necessary conditions for continuous parameter stochastic optimization problems*, SIAM J. Control **10** (1972), 550–565.

[LL1] J.-M. Lasry and P.-L. Lions, *Jeux à champ moyen. I. Le cas stationnaire*, C. R. Math. Acad. Sci. Paris **343** (2006), no. 9, 619–625 (French, with English and French summaries).

[LL2] J.-M. Lasry and P.-L. Lions, *Jeux à champ moyen. II. Horizon fini et contrôle optimal*, C. R. Math. Acad. Sci. Paris **343** (2006), no. 10, 679–684 (French, with English and French summaries).

[M] G. J. Mailath, *Do People Play Nash Equilibrium? Lessons From Evolutionary Game Theory*, Journal of Economic Literature **36** (1998), 1347–1374 (French, with English and French summaries).

[MS] J. Maynard Smith, *The theory of games and the evolution of animal conflicts*, J. Theoret. Biol. **47** (1974), no. 1, 209–221.

[MSP] J. Maynard Smith and G. R. Price, *The logic of animal conflict*, Nature **246** (1973), no. 5427, 15–18.

[M] R. B Myerson, *Nash equilibrium and the history of economic theory*, Journal of Economic Literature **37** (1999), no. 3, 1067–1082.

[N] J. Nash, *Non-cooperative games*, Ann. of Math. (2) **54** (1951), 286–295.

[vN] J. v. Neumann, *Zur Theorie der Gesellschaftsspiele*, Math. Ann. **100** (1928), no. 1, 295–320 (German).

[vNM] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, New Jersey, 1944.

[O] G. Owen, *Game theory*, 3rd ed., Academic Press, Inc., San Diego, CA, 1995.

[P] S. G. Peng, *A general stochastic maximum principle for optimal control problems*, SIAM J. Control Optim. **28** (1990), no. 4, 966–979.

[PP] H. J. Pesch and M. Plail, *The maximum principle of optimal control: a history of ingenious ideas and missed opportunities*, Control Cybernet. **38** (2009), no. 4A, 973–995.

[P] J.-P. Pier (ed.), *Development of mathematics 1950–2000*, Birkhäuser Verlag, Basel, 2000.

[S1] W. H. Sandholm, *Evolutionary Game Theory (Lecture Notes)*, Citeseer (2007).

[S2] W. H. Sandholm, *Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium*, Games **1** (2010), no. 1, 3–17.

[S3] W. H. Sandholm, *Stochastic evolutionary game dynamics: foundations, deterministic approximation, and equilibrium selection*, Evolutionary game dynamics, Proc. Sympos. Appl. Math., vol. 69, Amer. Math. Soc., Providence, RI, 2011, pp. 111–141.

[TJ] P. D. Taylor and L. B. Jonker, *Evolutionarily stable strategies and game dynamics*, Math. Biosci. **40** (1978), no. 1-2, 145–156.

[V] P. P. Varaiya, *On the existence of solutions to a differential game*, SIAM J. Control **5** (1967), 153–162.

[W] N. Williams, *On Dynamic Principal-Agent Problems in Continuous Time*, http://www.ssc.wisc.edu/∼nwilliam/ (Working Paper) (2013).

[YZ] J. Yong and X. Y. Zhou, *Stochastic Controls : Hamiltonian Systems and HJB Equations*, Stochastic Modelling and Applied Probability, Springer, New York, 1999.

[Ø] B. Øksendal, *Stochastic differential equations*, 6th ed., Universitext, Springer-Verlag, Berlin, 2003. An introduction with applications. MR2001996 (2004e:60102)