Games and information: An introduction to Game Theory

A book review by
Daniël Wedema

July 19, 2011

1 Overview

Title:	Games and information
	An introduction on Game Theory
Author:	Eric Rasmusen
First published:	1989
Type of book:	Educational reading/Science book
Subject:	Game Theory
Price:	EUR 39,99
ISBN:	978-1-4051-3666-2

2 Chapters

The book by Rasmusen [1] is divided in three main parts. Part I (chapters 1 to 6) gives a general overview of Game Theory. Part II (chapters 7 to 11) goes into detail about asymmetric information. And finally part III (chapters 12 to 14) gives examples about applications of the preceding theories.

2.1 Part I

The first part of the book discusses the basic terms used in Game Theory. Each chapter builds on the previous chapters. This also implies that a reader which is new to the subject matter can not go to a random chapter and fully understand it. It is necessary to read up on the previous chapters as the terms used in any chapter are usually already defined in the previous chapters.

2.1.1 Chapter 1

The first part of the book starts of in chapter 1 with introducing the rules of the game. Here the most important terms are defined which are used frequently throughout the book. Important definitions like 'players', 'actions', 'payoff' and 'information' are given. Although these definitions are clear now I have read the entire book, Rasmusen might have given a clearer connection between definition, notation and examples. The first chapter does use examples to more clearly define the used terms. So for example, the definition of payoff is given by $\pi_i(s_1, \cdots, s_n)$. Given the example of *The dry cleaners game*, it might

have been more clear when Rasmusen had specified the payoff in this game as $\pi_{\text{NewCleaner}}(\text{Enter}, \text{Low price}) = -100$ to make the connection between mathematical notation and the table with payoff values.

Furthermore in chapter 1 the concept of equilibrium is introduced. Connected to this subject is the process of finding equilibria. For this purpose the idea of dominant strategies and iterated dominance are introduced. Here again the connection between mathematical notation and the tabular values might have been made more explicit to make it easier for the reader to understand.

2.1.2 Chapter 2

Another important concept in Game Theory is the notion of Information. Chapter 2 deals with this subject. The main concept in this chapter is that of *Information set*, i.e. those states where a player can differentiate between. Of course information sets can have many properties, like having perfect, certain, symmetric or complete information. The idea of complete information however is ambiguous since it has an old definition and a new definition. Rasmusen has chosen to make this distinction clear in his book, and even gives a technique to transform one definition into the other. Although this is probably an essential move for someone familiar with the old definition and wanting to learn the new definition, it is confusing for a novice like me who has never seen any of the two definitions. For me it was not clear why I needed to know about the old definition and why I should be able to make the transformation to the new definition. This confusion is fueled specially by the fact that the old definition and the transformation is never seen again in the remaining parts of the book.

Another concept introduced in chapter 2 is that of beliefs, and the associated update rules using Bayes inference. Here a large part is reserved for writing out all the necessary calculations. What is missing however is the practical application of these calculations. What is the use of all these calculations?

2.1.3 Chapter 3

Chapters 1 and 2 only deal with discrete action spaces. Chapter 3 introduces mixed and continuous strategies. A mixed strategy is a strategy where the knowledge or believe of a player is mapped to a probability distribution over actions unlike the pure strategies where the believe of a player is mapped to a single action every time. There are multiple methods of finding equilibria when players use mixed strategies. Rasmusen introduced a first method which is based on calculus for maximizing the payoffs of the players.

After introducing this method, the book gives an interpretation of mixed strategies. Do these strategies actually exist? Does any player in the real world ever use a truly mixed strategy? However, Rasmusen does give a few nice possible interpretations of mixed strategy theory to justify this chapter. An example of such an interpretation might be a large group of people, each picking a pure strategy. However, when a random person is picked from this group, his behavior can be modeled using a mixed strategy.

After giving an interpretation of mixed strategies, Rasmusen continuous with a second method for finding equilibria. This ordering of paragraphs is a bit odd to me. It might have made more sense if he had started with introducing the concept of mixed strategy, and giving the interpretation right after that.

Once these ideas have been introduced it would follow naturally that he would continue with ways of finding equilibria. The second method of finding equilibria is called the 'payoff-equating' method, which is based on the idea that when a player uses a mixed strategy in equilibrium, he must be getting the same payoff from each of the pure strategies used in the mixed strategy. Otherwise the player would deviate to the pure strategy with a higher payoff. The idea behind this approach is different from the calculus based approach, but the actual calculations and the resulting solution are the same. After introducing these ideas, Rasmusen gives a few good examples to help secure the ideas.

The next paragraph generalizes the payoff-equating method to any number of players and general parameters. Here again I would have preferred a different ordering of subjects. It would have been better when the general method was introduced before the specific instantiation of this method to 2 players and specific parameters. This way this paragraph does not introduce really new concepts, instead of using actual numbers the calculations are now done with general parameters.

The next paragraph tries to explain the difference between randomized actions and a mixed strategy. Although Rasmusen does give an example to help interpret the difference between the two, it still is not really clear to me what the actual distinction is between the two. To me it still seems that randomized actions can always be transformed into a description using a mixed strategy.

The previous paragraphs were still using discrete action spaces, even though the strategy comes from a continuum. The next paragraph discusses what happens when the action space itself is a continuum. An example of this might be choosing a price for a product. Different methods for finding equilibria in these cases are the Stackelberg equilibrium and the Bertrand equilibrium. The difference between the two is in the order of actions. In a Stackelberg game one player picks a price after the other player. In a Bertrand game, both players pick a price at the same time.

Even though players might go for a mixed strategy, there does not always exist a equilibrium. It might for example be that the strategy space is unbounded, in which case each player might pick a strategy with an infinite value. Another possibility is that the strategy space is open, i.e. each player can pick a strategy in the range $[0, 100)$. In this case it might be that each of the players wants a strategy as close to 100 as possible. However, since the value 100 itself is out of range, the players will move closer and closer to 100 over time, but never reach an actual equilibrium. When the strategy space is discrete, it might also be possible that a number of players will not reach equilibrium. When the same game extends its rules to allow a continuous action space, it might suddenly be that an equilibrium does exist. A last possibility is when the reaction functions of two or more players simply do not intersect. The intersection point of these functions always give the equilibrium, so when this point does not exist a equilibrium does not exist as well.

A good example of this chapter and its connection to artificial intelligence can be found in [2]. Here the authors try to model agents who want to put a price on their products. The goal of these agents is to collect as much money as possible. Of course there are two factors that need to be optimized, i.e. the quantity of sold goods and the price of the sold goods. In their article they give a model of an automated market, something that does not seem very science fiction like at all. Such a market might very well be brought into existence some

day in the near future. They show in this article that the agents will move their price of the products very close to an equilibrium price. However, the equilibrium price is never actually reached because the prices they could pick from were discrete. The authors note that when a continuous price range would be accepted, the prices would drop to the Bertrand equilibrium. When they let the agents learn an optimal payoff function, they find that the automated learning gives consistent results. The optimum for the payoff function corresponds to the optimal price in an equilibrium. This implies that automated learning agents in the marketplace could actually be deployed and they would be profitable.

2.1.4 Chapter 4

In all previous chapters two players moved simultaneously. In real life however, players often move sequentially. In chapter 4 these situations are studied. Here the notion of Nash equilibrium is refined to a subgame perfect Nash equilibrium. Now not only the outcomes of the game need to be in a Nash equilibrium, but in every subgame that is encountered the strategy profile needs to be in a Nash equilibrium. This is relevant since there might be found many Nash equilibria, but not all of these are reasonable. Rasmusen illustrates this point nicely by the example of Nuisance Suits. In nuisance suits a plaintiff brings suit against a defendant. Usually the plaintiff knows that the probability of success when the case goes to trial is low. However, the costs for the defendant are also high, so the plaintiff hopes to settle the case outside of court and walk away with a positive payoff. The perfect equilibrium for these games can be found by backward reasoning.

2.1.5 Chapter 5

Chapter 5 discusses a subset of the sequential games called the repeated games. Here players move sequentially but in each round can choose an action from the same action set. The solution to such games is again done by backward inductive reasoning. For example in the repeated prisoner's dilemma one might think that it might be beneficial to build a reputation of being cooperative, i.e. *Deny* in each round. Once this 'trust' is established, the player can then choose *Confess* in the last round as the opponent will still choose *Deny*. However, both players will use this strategy, and thus each player will conclude that in the last round both will *Confess*. Because of this each player will conclude that the switch from *Deny* to *Confess* should be done in the next-to-last round. Repeating this backward induction shows that both players will actually *Confess* in all rounds.

This reasoning only succeeds for finitely repeated games since there is a base case to start the backward induction from. For infinitely repeated games finding the equilibrium becomes a bit harder. In this case a few strategies are discussed like the Grim strategy, Tit-for-tat and minimax strategies. Rasmusen also diverges from the economical and legal examples in this chapter by giving evolutionary example. This divergence gives a nice switch from the 'man-made' situations to a real life nature example of what strategies would be optimal.

2.1.6 Chapter 6

The previous chapters on dynamic games dealt only with complete information, i.e. all players know everything, and thus know the same. Real world situations

are of course different as each player has its own information, and not all is known. Chapter 6 deals with dynamic games with incomplete information. Revealing and hiding this information is an essential part of playing dynamic games. Finding equilibria in these situations is however a challenging task. The set of reasonable equilibria is again refined as was also done in the previous chapter. One of these refinements is found by defining the trembling-hand perfect equilibrium. This equilibrium is related to the perfect equilibrium of the previous chapter, with the exception that the strategy for any player should still be optimal when the other player mistakenly chooses an out-of-equilibrium action. Another refinement is based on the Bayesian equilibrium as discussed in chapter 2. This equilibrium is now refined to the perfect Bayesian equilibrium analogous to the perfect equilibrium defined in the previous chapter, but now for incomplete information and believes. Chapter 6 continues to give clear examples to illustrate the definitions of these new equilibria.

Another surprising fact is that a player his ignorance might work in his own advantage. For example when a player of intermediate strength does not know his own strength, it can commit to a relatively high effort, and discourage a weak opponent. When this same player commits to the same effort, it would like like a low effort from the perspective of a strong opponent and appease this strong player, giving itself a higher payoff.

In the next paragraph incomplete information is applied to the repeated prisoner's dilemma. Rasmusen describes the Gang of Four model, in which a few players are unable to play any strategy but Tit-for-Tat, and other players just pretend to be of that type. In this case, a small amount of incomplete information can make a big difference in the outcome. The Gang of Four theorem states that the number of times a player chooses *Confess* is independent on the total time the repeated prisoner's dilemma is played, but does depend on the probability of a player being a genuine Tit-for-Tat player.

Rasmusen goes on to describe the Axelrod tournament, in which a repeated prisoner's dilemma is played with computer agents. This way he experimentally collected data on which strategy would be the best. These experiments unequivocally showed that Tit-for-Tat was the strongest player, receiving the highest payoffs.

2.2 Part II

Part II of the book deals with asymmetric information. Asymmetry arises when one of the players has relevant information that (some of) the other players lack. This is such a broad topic with a lot of real world examples that Rasmusen devotes a few chapters of the book to this subject.

2.2.1 Chapter 7 and 8

These chapters deal with the subject of moral hazard. A distinction is made between two types of players; the principal (uninformed) player and the agent (informed) player. When these two players act together a principal-agent model is created. In this model the principal is usually a manager and the agent is a worker. The agent is payed a wage and has to put in some effort to his work. The payoff of the agent is increasing in the wage, and decreasing in the effort, i.e. the agent wants to do as little effort as possible for a maximum wage. The

principal on the other hand has a payoff which increases with the effort of the agent and decreases with the wage.

In this chapter Rasmusen introduces the production game and several variants thereof. This game is characterized by the principal who offers a contract with a wage to the agent, and the agent accepts or rejects the contract. When the agent accepts, he has to exert an effort as agreed upon in the contract. In the first version of the production game, every move is common knowledge, and the principal can calculate exactly what wage to offer the agent to do a certain amount of effort. Three types of contracts are discussed to cause the agent to exert the right amount of effort, e.g. a contract where the agent only gets paid when he exerts the right amount of effort, a contract where the agent only gets paid when he exerts at least an amount of effort or a wage which increases with the amount of effort.

In the second version of the game, it is not the principal who offers the contract, but it is the agent who makes the first move. In this case the information is still symmetric and the calculations are analogous to the first version of the game. The next versions of the game move away from the symmetric information to a situation with more uncertainty and imbalance in information. Rasmusen describes methods to find the right contract in each of these cases.

2.2.2 Chapter 9

Chapter 7 and 8 dealt with moral hazard. In these situations agents are identical. In chapter 9 games with adverse selection is introduced, in which agents are different. Because there are different types of agents, the principal can benefit from offering different types of contracts. Different types of agents are modeled in this book by letting each agent have a different kind of ability.

These models can also be used for example when goods are traded, since the situation can also be described as one where something is offered by one player and accepted or rejected by another. Rasmusen describes different types of sellers, and different tastes for the buyers. It is for example described what happens when the buyers all have identical tastes and there is a continuum of sellers. Another example is when the buyers value their products more than the sellers. A final example is given where the valuation of the sellers differ. The outcome in each of these examples can be very different, and Rasmusen does a good job of explaining what will happen in each of these situations.

The chapter finishes with a variety of applications, e.g. in health insurance and bank loans.

2.2.3 Chapter 10 and 11

Chapters 10 and 11 both deal with information exchange between agents. In chapter 10 this is done using a mechanism design, which is a set of rules that one player constructs and another player accepts. The mechanism that is transferred contains an information report by the accepting player, and the player that made the mechanism design uses this to select his actions. An example of this could be an insurance company which has two types of insurances; one for the high risk agents and one for the low risk agents. In the information report the agents can communicate what type of agent they are, and the insurance company can then offer the correct insurance.

Of course a problem here is when the agent is able to lie about the information it provides, e.g. a high risk agent telling the insurance company that it is a low risk agent in order to get a cheaper insurance. The book mentions a few ways how to force the agent to tell the truth. An example could be by implementing a penalty for lying whenever the insurance company discovers that the agent has lied.

Chapter 11 deals with another form of information exchange, namely signaling. The signal is something the agent chooses for himself and on which the principal can base a wage. This notion is again introduced by Rasmusen with a number of examples in the context of education to make it more clear.

At the end of chapter 11 a few ways to communicate are mentioned, among which cheap talk, auditing, mechanism, signaling and screening. This summary of the different kinds of communication are nice to have since these subjects were discussed throughout the book, but are important enough to have in one short list.

2.3 Part III

In part I and part II of the book the majority of game theory subjects are covered. To make this theory better understandable, Rasmusen has filled these parts with numerous examples. Part III extends this list of examples considerably by giving real world applications of the theory discussed in the book.

This part is also the only one where the chapters can be read in any order without lacking any information. In part I and part II this was not possible since all chapters were build upon the preceding chapters. However, to read any of the chapters of part III, it is necessary to read parts I and II first.

2.3.1 Chapter 12

Chapter 12 deals with the subject of bargaining. When many agents act in a market to determine the equilibrium price, standard economic theories apply. For example the intersection of the supply and demand curve can be found to find the equilibrium price and quantity for goods. However, when there are a small number of agents in some market, usually the price is not set, but a bargaining process will take place. In previous chapters this situation was not encountered, since the agent only had the choice to accept or reject an offer.

An example of a bargaining game which is introduced in this chapter is that of Splitting a pie. Two players must decide on a share of a pie they want for themselves. When the shares of both players combined is greater than the pie, both get nothing otherwise each gets the share they requested.

Nash proposed a solution to the bargaining problem in 1950. To formulate a solution he had to introduce four axioms which had to be fulfilled. If any solution satisfies the four axioms, then it is a unique optimal strategy profile.

In the original splitting a pie example, both players can make a single action and they do this simultaneously. A simple adaptation to this would be when both players move sequentially and can observe each other his offers. When a player makes an offer, the other player can either accept or reject and make a counteroffer. The book mentions a method using backward inductive reasoning as was also used in chapter 5.

Backward inductive reasoning becomes more difficult when both players go on for eternity with rejecting and making counter offers. However, the method used in a finite game can be generalized to also calculate the equilibrium in the infinite game.

2.3.2 Chapter 13

Chapter 13 gives an extension of chapter 12. Where chapter 12 discusses bargaining (a situation with one seller and one buyer), chapter 13 discusses auctions. At an auction there is a single seller and multiple buyers. Auctions are usually more formal and rigid, and are thus better suited to be modeled by game theory.

Different examples are presented by Rasmusen, e.g. when the price to be determined is on a discrete scale versus when the price is on a continuous scale. In the next paragraph Rasmusen describes optimal strategies in these different versions.

2.3.3 Chapter 14

In chapter 14 the Cournot equilibrium as originally described in chapter 3 is revisited. It goes more into detail about how firms in a market set their prices. The Cournot game is slightly extended by letting the two firms choose the quantity they want to sell.

3 Closing remarks

The book Games and Information: An introduction to game theory by Eric Rasmusen is a thorough description of the current state of game theory. Although many subjects were treated, I was missing a good overview of how all subjects fitted together. In my opinion it would have been better when at the beginning of chapters Rasmusen would have told us what was going to be told and how it all fits together. Rasmusen does give a short overview at the beginning of each chapter, but in doing this uses terms which are explained in the succeeding chapter. Even though this is understandable, it causes the reader to only understand the introduction after completing the chapter, at which point the introduction is not very useful anymore.

Rasmusen has made a good effort at balancing the amount of mathematical notation and examples by stories. However, sometimes this does feel a bit awkward, e.g. when in an example of a game the players have concrete names, but the payoffs are still abstract numbers without any direct link to reality.

Rasmusen is a theoretical economist, which can be clearly seen by the examples that he comes up with. These examples are strictly tied to the real world which might make it easier for people to understand. However, for someone like me without any economic background, these example sometimes make it hard to immediately understand what is going on. In those cases I have to master the new game theoretical concepts as well as new economic concepts. For me it might have been easier to understand when examples would have come from actual games (e.g. board games, card games, etc.). When these subject would have been used in parts I and II of the book, the real economical examples could have been saved for part III. However, this critique only applies to students having no economical background, which is not who Rasmusen had in mind while

writing this book. After all, this book is used by Rasmusen to teach his students at the School of Business in Indiana, so it makes sense that it is aimed at such an audience.

Luckily the book is self-contained. Rasmusen does give plenty of pointers to articles which can be used when someone wants to know more in a specific subject. However, for a general overview and introduction of the field of game theory, only his book would suffice. For this reason I would recommend this book to my fellow students from the field of AI or any related field.

References

[1] E. Rasmusen. *Games and information: An introduction to game theory.* Blacwell publishing, fourth edition, 2007.

[2] E. Kutschinski, T. Uthmann, and D. Polani. Learning competitive pricing strategies by multi-agent reinforcement learning. *Journal of Economic Dynamics and Control*, 27(11-12):2207–2218, 2003.

We realize that games are first and foremost about fun and our greatest ambition is to create only the most enjoyable and amusing playware. The KDEGames team achieves this goal by paying greater attention to gaming details, using the latest KDE based technology, and listening closely to our users' feedback. To date KDEGames has accumulated an impressive selection of quality games in nearly all genres. Each of these small masterpieces have been carefully tailored to satisfy the