DOWNLOAD PDF V. 2. RATIONALITY, GAMES, AND STRATEGIC BEHAVIOUR.

Chapter 1 : Price Theory, Chapter Game Theory, Strategic Behavior, and Oligopoly

Get this from a library! Collected papers in theoretical economics. Volume II, Rationality, games, and strategic behaviour. [Kaushik Basu].

 The well-known Rational Pigs matrix game has been relatively intensively analyzed in terms of reassessment of the logic of two players involved in asymmetric situations as gluttons that differ significantly by their attributes. This paper presents a successful attempt of using autogenerator for creating the framework of the game, including the predefined scenarios and corresponding payoffs. Autogenerator offers flexibility concerning the specification of game parameters, which consist of variations in the number of simultaneous players and their features and game objects and their attributes as well as some general game characteristics. In the proposed approach the model of autogenerator was upgraded so as to enable program specification updates. For the purpose of treatment of more complex strategic scenarios, we created the Rational Pigs Game Extended RPGE , in which the introduction of a third glutton entails significant structural changes. Introduction This paper discusses the issue of simulation of strategic games that fall within the domain of game theory analysis. As a discipline of interactive decision-making, game theory deals with noncooperative as well as cooperative scenarios. However, game theory is primarily concerned with analyzing competitions or conflict situations, defining the choices for each individual player and examining possible resulting outcomes and behaviors in such competitive games [1]. Deeper insights into the established game theory scenarios or models have been enabled by the application of computer programs. Yet tractability limitations have been found to be a common disadvantage of computer-performed simulation models [2]. Namely, the decision process is hard to trace even if the source code is provided. The modelling framework for the development of strategic games scenarios described in this paper gives a detailed view of the tool on which the simulation is performed and thus increases the tractability of simulations. A literature review reveals that common approaches in the field of game theory are the use of agents [3] and different models of learning such as the Monte Carlo method and temporal difference learning [4]. The intention of our approach entirely focused on real players is to provide a modelling framework that would be easily upgradable with new features. One of the gluttons is relatively large, which brings asymmetry into the strategic situation. The food is released into the cage at the opposite side of the lever, which presents a handicap for the glutton who actually presses the lever because of the delay in his feeding activity. Although rational decision-making in the afore-described situation has been analyzed by a number of authors such as [6 â€” 8] and others, the most acknowledged analysis is generally considered to be the one by McMillan [9]. Gintis in [10] discusses a game which involves a large and a small monkey, where the strategic framework is almost identical since it is necessary for the provision of food to climb a tree with fruits and shake the tree. According to Czarnecki and Eisenecker [11], generative programming is a discipline within automatic programming aimed at automating the software development process. In general, such automatically produced software is stored for later use e. On the other hand, in the autogenerator-based approach [12] the necessary piece of code is produced and immediately executed on demand. Instead of program files, in the case of autogenerator the generated code is stored into variables and also executed from them. The concept also uses the convenient capability of scripting languages to autoevaluate the programming code. In our particular case Python is used owing to the flexibility of its supported data structures e. In the simulation of the elaborated strategic game, program specification is used for code generation as well as for recoding a set of game properties in course of the game. Such properties can refer to the game as a whole like the state of a semaphore and the game phase and tanks their position, amount of fuel, open or closed state, etc. The first issue examined concerns the stability of game solutions considering two gluttons in a new situation that includes a third player. Modelling a Framework for the Development of

Games of Strategy 2. Introduction to Autogenerator Autogenerator is a model of application development where program code is generated and executed on demand. It is based on the SCT generator model [13] that is aimed to produce entire applications instead of skeletons that require additional work. Autogenerator represents an example of the advanced use of frame-based software development. This increases flexibility in the development of generators with regard to the use of static frames. Configuration is a set of rules that manage the generation process, where programming code is produced by assembling the features from Specification, together with code artifacts stored in Templates. Templates refer to a set of code artifacts that are used as building blocks of generated applications. Specification, Configuration, and Templates together build an SCT frame, which contains all the information needed by the SCT generator to produce program code. The autogeneration process is shown in Figure 1. The user sends a request which contains information about the user and the action that the user wants to take. The request is accepted by the request handler, whose task is to decompose the request, determine what action to take, and call the source code generator to produce the appropriate source code. The generated source code is then stored in a variable, where it can be evaluated by scripting languages like JavaScript, Perl, or Python. The generated source code is evaluated by the execution unit, as shown in Figure 1. The execution unit executes the generated source code together with the arguments provided by the request handler. Those arguments are represented in Figure 1 as the application context. The execution unit sends the result to the user. In the former version of autogenerator three key new features of generated software were introduced. The second one is the usage of imperative instructions in Specification. Such instructions are aimed to be performed only once, usually to harmonize the program code with some program dependency e. After they have been executed, imperative instructions are deleted from Specification. The third feature is introspection which, in the case of autogenerator, gives the developer an insight into basic model elements that are used in the production of a specific piece of code e. The model of autogenerator was extended for the purpose of implementing a simulation mechanism, like the strategic game in our example, by introducing the possibility of changing Specification from the autogenerated application. This includes operations such as changing attribute values, adding new attributes, deleting the attributes, and testing their values. Each change in Specification changes the generated code, making the entire application dynamic. All the aforementioned features are applied in the program example. Autogenerator in Strategic Games The model of the usage of autogenerator in strategic games is presented in Figure 2. The implementation of the game is performed by means of configuration and templates of the program code. Each strategic game has its own configuration and its own set of program code templates. The configuration defines the way in which program code templates are combined to produce the final program code. If a new feature is to be introduced into a strategic game, new templates will have to be created and the configuration will have to be updated. The introduction of new features is possible even if a strategic game is running since autogenerator produces program code on demand. The possibility of changing the rules in a game that is running is thus an important feature of using autogenerator in strategic games. Model of using autogenerator in strategic games. The number of players and their parameters can change over time since they represent real-life scenarios. The number of players and their properties are defined in Specification by a set of parameters. Each player has their own set of properties that can be updated, extended, or reduced during the game. Since some strategic games can run longer, adding new players and changing their properties during the game make for another key feature of autogenerator usage in the context of strategic games. The implementation of the presented model of autogenerator in strategic games is described in detail on the model of the Rational Pigs game discussed in Section 4. Extended Autogenerator Model Although the feature of regenerating program code on demand was already included in the original model of autogenerator presented in [12], the possibility of changing its Specification by the autogenerated application in form of imperative instructions, as described below was limited. To fulfill the requirements of different simulations, it would be useful to specify the attributes of simulated objects in Specification and changing values of its attributes during execution. The updated attributes and their values could be used in the following regeneration cycle in production and execution of the new program code Figure 2. Such frequent reading and updating of Specification may give rise to some implementation issues, especially in the case of concurrent access by several autogenerated processes, as discussed in Section 3. On the other hand, the possibility of updating Specification can eliminate the need for usage of any external data sources within simulations like strategic games. The demand can be sent by the user e. Updating attributes and their values. It is also possible to check the existence of an attribute and its value. There are five updating functions in the current model Table 1. By using updating functions, Specification also assumes the role of a small database that can be used in further generation of the code and the report that contains the values obtained during execution. Design and Implementation Issues Two main issues occurred during the process of designing updating functions. The first one concerns addressing the attributes in Specification. Although Specification has a tree structure, which suggests the usage of paths similarly to addressing folders on a hard disk, such paths can be long and should be changed in case of tree restructuring. In the approach presented in this paper, names and values of parent attributes are therefore used instead of paths, with the assumption that their names are unique. Taking into account this limitation, it is simple to address parts of Specification as in the following example: In some cases, Specification was reached by a process at the moment when it was only partially updated, which led to a temporal collapse of the autogenerated application. Each updating function deals with Specification in the same way: As shown in Figure 4 , each player accesses the dynamic application which is produced by autogenerator on demand. This demand could be given directly by the player e. Extended autogenerator model in a strategic game example. During the game, the values of the attributes in Specification change, reflecting the actual state. In addition, values achieved during the simulation can be easily read from Specification. Consideration of rationality in this game has its roots in the famous experiment from , reported by Baldwin and Messe, wherein the intelligence of animals was tested [16]. On the example of the strategic game in Table 2 , we analyze the decisions of participants and search for an equilibrium state, that is, the solution to the game. Boxed Pigs model, as in [9]. While some authors suggest payoffs that are somewhat different from those shown above, making generalizations as in [17], others reduce the importance of the effort and the associated cost due to the pressing of the lever, as discussed in [6], which practically leads to possible negative payoff for Sp. Therefore, in our example the amounts of the payoffs matrix are somewhat loosely defined as The solution to the presented game is calculated using the procedure of Nash equilibrium, using arguments proposed in [18]. The utility that a player i realizes in the game is denoted by. The space of pure strategies of player i is denoted by. Important theoretical and behavioural questions about the rationality of players emerge in practical analogies of such scenarios. Analysts such as McMillan [9] examined the logic of Dp through the analogy with strategic scenarios of a real-life cartel OPEC , where Saudi Arabia assumed the role of the dominant glutton. In the repetitive scenario of this game, the existence of mixed Nash equilibria is possible, with the coresponding rationality of choosing probabilities frequencies of strategies. We anticipate this game as a one-stage game.

DOWNLOAD PDF V. 2. RATIONALITY, GAMES, AND STRATEGIC BEHAVIOUR.

Chapter 2 : Decision Making Models: Rational and Behaviour Model

Request PDF on ResearchGate | Strategic behaviour, rationality and bargaining power in two-stage games with artificial socially inspired agents | In this paper we demonstrate that artificial.

Johnny Von Neumann and the rest of us. An economy is an interdependent system. In the process of solving it we have deliberately pushed that interdependency into the background. The rest of the world consists for him of a set of prices--prices at which he can sell what he produces and buy what he wants. The monopolist of Chapter 10 is big enough to affect the entire market, but he is dealing with a multitude of individual consumers. From the standpoint of the monopolist, the customer is not a person at all; he is simply a demand curve. Our analysis has thus eliminated an important feature of human interaction and of many markets--bargaining, threats, bluffs, the whole gamut of strategic behavior. That is one of the reasons why most of price theory seems, to many students, such a bloodless abstraction. We are used to seeing human society as a clash of wills, whether in the boardroom, on the battlefield, or in our favorite soap opera. Economics presents it instead in terms of solitary individuals, or at the most small teams of producers, each calmly maximizing against an essentially nonhuman environment, an opportunity set rather than a population of self-willed human beings. There is a reason for doing economics this way. The analysis of strategic behavior is an extraordinarily difficult problem. John Von Neumann, arguably one of the smartest men of this century, created a whole new branch of mathematics in the process of failing to solve it. The work of his successors, while often ingenious and mathematically sophisticated, has not brought us much closer to being able to say what people will or should do in such situations. Seen from one side, what is striking about price theory is the unrealistic picture it presents of the world around us. Seen from the other, one of its most impressive accomplishments is to explain a considerable part of what is going on in real markets while avoiding, with considerable ingenuity, any situation involving strategic behavior. When it fails to do so, as in the analysis of oligopoly or bilateral monopoly, it rapidly degenerates from a coherent theory to a set of educated guesses. What Von Neumann created, and what this chapter attempts to explain, is game theory. I start, in Part 1, with an informal description of a number of games, designed to give you a feel for the problems of strategic behavior. Part 2 contains a more formal analysis, discussing various senses in which one might "solve" a game and applying the solution concepts to a number of interesting games. Parts 3 and 4 show how one can attempt, with limited success, to apply the ideas of game theory to specific economic problems. Strategic Behavior "Scissors, Paper, Stone" is a simple game played by children. At the count of three, the two players simultaneously put out their hands in one of three positions: The winner is determined by a simple rule: The game may be represented by a 3x3 payoff matrix, as shown in Figure Rows represent strategies for player 1, columns represent strategies for Player 2. Each cell in the matrix is the intersection of a row and a column, showing what happens if the players choose those two strategies; the first number in the cell is the payoff to Player 1, the second the payoff to Player 2. It is convenient to think of all payoffs as representing sums of money, and to assume that the players are simply trying to maximize their expected return--the average amount they win--although, as you will see, game theory can be and is used to analyze games with other sorts of payoffs. Figure The payoff matrix for Scissors, Paper, Stone. The top left cell shows what happens if both players choose scissors; neither wins, so the payoff is zero to each. The next cell down shows what happens if Player 1 chooses paper and Player 2 chooses scissors. Scissors cuts paper, so Player 2 wins and Player 1 loses, represented by a gain of one for Player 2 and a loss of one for Player 1. I have started with this game for two reasons. The first is that, because each player makes one move and the moves are revealed simultaneously, it is easily represented by a matrix such as Figure , with one player choosing a row, the other choosing a column, and the outcome determined by their intersection. We will see later that this turns out to be a way in which any two-person game can be represented, even a complicated one such as chess. The second reason is that although this is a simple game, it is far from clear what its solution is--or even what it means to solve it. Some quite complicated games have a winning strategy for one of the players. But there is no such strategy for Scissors, Paper, Stone. Whatever you choose is right or wrong only in relation to what the other player chooses. While it may be hard to say what the correct strategy is, one can say with some confidence that a player who always chooses stone is making a mistake; he will soon find that his stone is always covered. One feature of a successful strategy is unpredictability. That insight suggests the possibility of a deliberately randomized strategy. Suppose I choose my strategy by rolling a die, making sure the other player is not watching. If it comes up 1 or 2, I play scissors; 3 or 4, paper; 5 or 6, stone. Whatever strategy the other player follows other than peeking at the die or reading my mind , I will on average win one third of the games, lose one third of the games, and draw one third of the games. Can there be a strategy that consistently does better? Not against an intelligent opponent. The game is a symmetrical one; the randomized strategy is available to him as well as to me. If he follows it then, whatever I do, he will on average break even, and so will I. One important feature of Scissors, Paper, Stone is that it is a zero-sum game; whatever one player wins the other player loses. While there may be strategy of a sort in figuring out what the other player is going to do, much of what we associate with strategic behavior is irrelevant. There is no point in threatening to play stone if the opponent does not agree to play scissors; the opponent will refuse, play paper, and cover your stone. Bilateral Monopoly, Nuclear Doom, and Barroom Brawls Consider next a game discussed in an earlier chapter--bilateral monopoly. The rules are simple. You and I have a dollar to divide between us, provided that we can agree on a division. If we cannot agree, the dollar vanishes. This game is called bilateral monopoly because it corresponds to a market with one buyer and one seller. The apple is worth nothing to me and one dollar to you. If I sell it to you for a dollar, I am better off by a dollar and you, having paid exactly what the apple is worth, are just as well off as if you had not bought it. If I give it to you, I gain nothing and you gain a dollar. Any price between one and zero represents some division of the dollar gain between us. If we cannot agree on a price I keep the apple and the potential gain from the trade is lost. Bilateral monopoly nicely encapsulates the combination of common interest and conflict of interest, cooperation and competition, typical of many human interactions. The players have a common interest in reaching agreement but a conflict over what the terms of the agreement will be. The United States and the Soviet Union have a common interest in preserving peace but a conflict over how favorable the terms of that peace will be to each side. Husband and wife have a common interest in preserving a happy and harmonious marriage but innumerable conflicts over how their limited resources are to be spent on things that each values. Members of a cartel have a common interest in keeping output down and prices up but a conflict over which firm gets how much of the resulting monopoly profit. Bilateral monopoly is not a zero-sum game. That makes it fundamentally different from Scissors, Paper, Stone; it permits threats, bargains, negotiation, bluff. I decide to get 90 cents of the dollar gain. I inform you that I will refuse to accept any less favorable terms; you may choose between 10 cents and nothing. If you believe me, you give in. If you call my bluff and insist that you will only give me 40 cents, I in turn, if I believe you, have the choice of 40 cents or nothing. Each player is trying to get a better outcome for himself by threatening to force an outcome that is worse for both. One way to win such a game is to find some way to commit oneself, to make it impossible to back down. A child with good strategic instincts might announce "I promise not to let you have more than 20 cents of the dollar, cross my heart and hope to die. The second player goes home with 20 cents and a resolution that next time he will get his promise out first. The strategy of commitment is not limited to children. Its most dramatic embodiment is the doomsday machine, an idea dreamed up by Hermann Kahn and later dramatized in the movie Doctor Strangelove. Suppose the United States decides to end all worries about Soviet aggression once and for all. It does so by building a hundred cobalt bombs, burying them in the Rocky Mountains, and attaching a fancy geiger counter. If they go off, the cobalt bombs produce enough fallout to eliminate all human life anywhere on earth. The geiger counter is the trigger, set to explode the bombs if it senses the radiation from a Soviet attack. We can now dismantle all other defenses against nuclear attack; we have the ultimate deterrent. In an improved version, dubbed by Kahn the Doomsday-in-a-hurry Machine, the triggering device is somehow equipped to detect a wide range of activities and respond accordingly; it could be programmed, for instance, to blow up the world if the Soviets invade West Berlin, or West Germany, or anywhere at all--thus saving us the cost of a conventional as well as a nuclear defense. While a doomsday machine is an elegant idea, it has certain problems. In Doctor Strangelove, it is the Russians who build one. Unfortunately, while they are waiting, a lunatic American air force officer launches a nuclear strike against the Soviet Union. The doomsday machine is not entirely imaginary. Consider the situation immediately after the United States detects the beginning of an all-out nuclear strike by the Soviet Union. Assume that, as is currently the case, we have no defenses, merely the ability to retaliate. The threat of retaliation may prevent an attack, but if the attack comes anyway retaliation will not protect anyone. It may even, by increasing fallout, climactic effects, and the like, kill some Americans--as well as millions of Russians and a considerable number of neutrals who have the misfortune to be downwind of targets. Retaliation in such a situation is irrational. Nonetheless, it would probably occur. The people controlling the relevant buttons--bomber pilots, air force officers in missile silos, nuclear submarine captains--have been trained to obey orders. They are particularly unlikely to disobey the order to retaliate against an enemy who has just killed, or is about to kill, most of their friends and family. Our present system of defense by retaliation is a doomsday machine, with human beings rather than geiger counters as the trigger. So far both have worked, with the result that neither has been used.

DOWNLOAD PDF V. 2. RATIONALITY, GAMES, AND STRATEGIC BEHAVIOUR.

Chapter 3 : Evolutionary game theory - Wikipedia

(v) The problem requiring a solution may be complex and unstructured, it may not be defined with rationality. In such situations a manager may rely on intuition than on rationality, (vi) The organizational variables such as philosophy, multiplicity of goals, existence of informal goals, power structure may be taken into account while taking.

The game is simple â€” rock beats scissors blunts it , scissors beats paper cuts it , and paper beats rock wraps it up. Anyone who has ever played this simple game knows that it is not sensible to have any favoured play â€” the opponent will soon notice this and switch to the winning counter-play. The best strategy a Nash equilibrium is to play a mixed random game with any of the three plays taken a third of the time. This, in EGT terms, is a mixed strategy. But many lifeforms are incapable of mixed behavior â€" they only exhibit one strategy known as a pure strategy. If the game is played only with the pure Rock, Paper and Scissors strategies the evolutionary game is dynamically unstable: Rock mutants can enter an all scissor population, but then â€" Paper mutants can take over an all Rock population, but then â€" Scissor mutants can take over an all Paper population â€" and on and on This is easily seen on the game payoff matrix, where if the paths of mutant invasion are noted, it can be seen that the mutant "invasion paths" form into a loop. This in triggers a cyclic invasion pattern. A computer simulation of the Rock Scissors Paper game. Starting with an arbitrary population the percentage of the three morphs builds up into a continuously cycling pattern. Rock-paper-scissors incorporated into an evolutionary game has been used for modelling natural processes in the study of ecology. The social cyclic behaviors, predicted by evolutionary game theory, have been observed in various laboratory experiments. The overall situation corresponds to the Rock, Scissors, Paper game, creating a six-year population cycle. When he read that these lizards were essentially engaged in a game with rock-paper-scissors structure, John Maynard Smith is said to have exclaimed "They have read my book! Signalling theory Aside from the difficulty of explaining how altruism exists in many evolved organisms, Darwin was also bothered by a second conundrum â€" why do a significant number of species have phenotypical attributes that are patently disadvantageous to them with respect to their survival â€" and should by the process of natural section be selected against â€" e. On analysis, problems of biological life are not at all unlike the problems that define economics â€" eating akin to resource acquisition and management , survival competitive strategy and reproduction investment, risk and return. Game theory was originally conceived as a mathematical analysis of economic processes and indeed this is why it has proven so useful in explaining so many biological behaviours. A simple model of cost assumes that all competitors suffer the same penalty imposed by the Game costs, but this is not the case. More successful players will be endowed with or will have accumulated a higher "wealth reserve" or "affordability" than less successful players. This wealth effect in evolutionary game theory is represented mathematically by " resource holding potential RHP " and shows that the effective cost to a competitor with higher RHP are not as great as for a competitor with a lower RHP. As a higher RHP individual is more desirable mate in producing potentially successful offspring, it is only logical that with sexual selection RHP should have evolved to be signalled in some way by the competing rivals, and for this to work this signalling must be done honestly. Amotz Zahavi has developed this thinking in what is known as the handicap principle , [48] where superior competitors signal their superiority by a costly display. As higher RHP individuals can properly afford such a costly display this signalling is inherently honest, and can be taken as such by the signal receiver. Nowhere in nature is this better illustrated than in the magnificent and costly plumage of the peacock. The mathematical proof of the handicap principle was developed by Alan Grafen using evolutionary game-theoretic modelling. Evolutionary games which lead to a stable situation or point of stasis for contending strategies which result in an evolutionarily stable strategy Evolutionary games which exhibit a cyclic behaviour as with RPS game where the proportions of contending strategies continuously cycle over time within the overall population Competitive Co-evolution - The rough-skinned newt Tarricha granulosa is highly toxic, due to an evolutionary arms race with a predator, the common garter snake Thamnophis sirtalis , which in turn is highly tolerant of the poison. The two are locked in a Red Queen arms race. A third, co-evolutionary , dynamic combines intra-specific and inter-specific competition. Examples include predator-prey competition and host-parasite co-evolution, as well as mutualism. Evolutionary game models have been created for pairwise and multi-species coevolutionary systems. In competitive non-mutualistic inter-species coevolutionary system the species are involved in an arms race â€" where adaptations that are better at competing against the other species tend to be preserved. Both game payoffs and replicator dynamics reflect this. This leads to a Red Queen dynamic where the protagonists must "run as fast as they can to just stay in one place". A key factor applicable in these coevolutionary systems is the continuous adaptation of strategy in such arms races. Coevolutionary modelling therefore often includes genetic algorithms to reflect mutational effects, while computers simulate the dynamics of the overall coevolutionary game. The resulting dynamics are studied as various parameters are modified. Because several variables are simultaneously at play, solutions become the province of multi-variable optimisation. The mathematical criteria of determining stable points are Pareto efficiency and Pareto dominance, a measure of solution optimality peaks in multivariable systems. Darwinian assumptions about fitness are modeled using replicator dynamics to show that the organism evolving at a slower rate in a mutualistic relationship gains a disproportionately high share of the benefits or payoffs. With this understanding in place it is then appropriate to see if other, more subtle, parameters second order effects further impact the primary behaviours or shape additional behaviours in the system. Some of these key extensions to EGC are: A Spatial Game In a spatial evolutionary game contestants meet in contests at fixed grid positions and only interact with immediate neighbors. Shown here are the dynamics of a Hawk Dove contest, showing Hawk and Dove contestants as well as the changes of strategy taking place in the various cells Spatial Games Geographic factors in evolution include gene flow and horizontal gene transfer. Spatial game models represent geometry by putting contestants in a lattice of cells: Winning strategies take over these immediate neighbourhoods and then interact with adjacent neighbourhoods. Spatial structure is sometimes abstracted into a general network of interactions. This models the reality of most normal social interactions which are non-kin related. Unless a probability measure of reputation is available in Prisoners Dilemma only direct reciprocity can be achieved. Alternatively, agents might have access to an arbitrary signal initially uncorrelated to strategy but becomes correlated due to evolutionary dynamics. This is the green-beard effect or evolution of ethnocentrism in humans.

Chapter 4 : Game Theory (Stanford Encyclopedia of Philosophy)

The rise of game theory around the middle of the twentieth century has profoundly affected the way economic theory is practiced. The present book is a collection of Kaushik Basu s papers on game theory and, more generally, strategic analysis.

For reasons to be discussed later, limitations in their mathematical framework initially made the theory applicable only under special and limited conditions. This situation has dramatically changed, in ways we will examine as we go along, over the past six decades, as the framework has been deepened and generalized. Refinements are still being made, and we will review a few outstanding problems that lie along the advancing front edge of these developments towards the end of the article. Despite the fact that game theory has been rendered mathematically and logically systematic only since , game-theoretic insights can be found among commentators going back to ancient times. Consider a soldier at the front, waiting with his comrades to repulse an enemy attack. But if he stays, he runs the risk of being killed or woundedâ€"apparently for no point. On the other hand, if the enemy is going to win the battle, then his chances of death or injury are higher still, and now quite clearly to no point, since the line will be overwhelmed anyway. Based on this reasoning, it would appear that the soldier is better off running away regardless of who is going to win the battle. Of course, this point, since it has occurred to us as analysts, can occur to the soldiers too. Does this give them a reason for staying at their posts? If each soldier anticipates this sort of reasoning on the part of the others, all will quickly reason themselves into a panic, and their horrified commander will have a rout on his hands before the enemy has fired a shot. Long before game theory had come along to show analysts how to think about this sort of problem systematically, it had occurred to some actual military leaders and influenced their strategies. Thus the Spanish conqueror Cortez, when landing in Mexico with a small force who had good reason to fear their capacity to repel attack from the far more numerous Aztecs, removed the risk that his troops might think their way into a retreat by burning the ships on which they had landed. With retreat having thus been rendered physically impossible, the Spanish soldiers had no better course of action but to stand and fightâ€"and, furthermore, to fight with as much determination as they could muster. He took care to burn his ships very visibly, so that the Aztecs would be sure to see what he had done. They then reasoned as follows: Any commander who could be so confident as to willfully destroy his own option to be prudent if the battle went badly for him must have good reasons for such extreme optimism. The Aztecs therefore retreated into the surrounding hills, and Cortez had his victory bloodlessly. These two situations, at Delium and as manipulated by Cortez, have a common and interesting underlying logic. Notice that the soldiers are not motivated to retreat just, or even mainly, by their rational assessment of the dangers of battle and by their self-interest. Rather, they discover a sound reason to run away by realizing that what it makes sense for them to do depends on what it will make sense for others to do, and that all of the others can notice this too. Even a quite brave soldier may prefer to run rather than heroically, but pointlessly, die trying to stem the oncoming tide all by himself. Thus we could imagine, without contradiction, a circumstance in which an army, all of whose members are brave, flees at top speed before the enemy makes a move. What we have here, then, is a case in which the interaction of many individually rational decision-making processesâ€"one process per soldierâ€"produces an outcome intended by no one. Most armies try to avoid this problem just as Cortez did. During the Battle of Agincourt Henry decided to slaughter his French prisoners, in full view of the enemy and to the surprise of his subordinates, who describe the action as being out of moral character. The reasons Henry gives allude to non-strategic considerations: However, a game theorist might have furnished him with supplementary strategic and similarly prudential, though perhaps not moral justification. His own troops observe that the prisoners have been killed, and observe that the enemy has observed this. Metaphorically, but very effectively, their boats have been burnt. The slaughter of the prisoners plausibly sent a signal to the soldiers of both sides, thereby changing their incentives in ways that favoured English prospects for victory.

These examples might seem to be relevant only for those who find themselves in sordid situations of cut-throat competition. Perhaps, one might think, it is important for generals, politicians, mafiosi, sports coaches and others whose jobs involve strategic manipulation of others, but the philosopher should only deplore its amorality. Such a conclusion would be highly premature, however. The study of the logic that governs the interrelationships amongst incentives, strategic interactions and outcomes has been fundamental in modern political philosophy, since centuries before anyone had an explicit name for this sort of logic. Philosophers share with social scientists the need to be able to represent and systematically model not only what they think people normatively ought to do, but what they often actually do in interactive situations. The best situation for all people is one in which each is free to do as she pleases. Often, such free people will wish to cooperate with one another in order to carry out projects that would be impossible for an individual acting alone. But if there are any immoral or amoral agents around, they will notice that their interests might at least sometimes be best served by getting the benefits from cooperation and not returning them. Suppose, for example, that you agree to help me build my house in return for my promise to help you build yours. After my house is finished, I can make your labour free to me simply by reneging on my promise. I then realize, however, that if this leaves you with no house, you will have an incentive to take mine. This will put me in constant fear of you, and force me to spend valuable time and resources guarding myself against you. I can best minimize these costs by striking first and killing you at the first opportunity. Of course, you can anticipate all of this reasoning by me, and so have good reason to try to beat me to the punch. Since I can anticipate this reasoning by you, my original fear of you was not paranoid; nor was yours of me. In fact, neither of us actually needs to be immoral to get this chain of mutual reasoning going; we need only think that there is some possibility that the other might try to cheat on bargains. Once a small wedge of doubt enters any one mind, the incentive induced by fear of the consequences of being preemptedâ€”hit before hitting firstâ€”quickly becomes overwhelming on both sides. If either of us has any resources of our own that the other might want, this murderous logic will take hold long before we are so silly as to imagine that we could ever actually get as far as making deals to help one another build houses in the first place. The people can hire an agentâ€”a governmentâ€—whose job is to punish anyone who breaks any promise. So long as the threatened punishment is sufficiently dire then the cost of reneging on promises will exceed the cost of keeping them. The logic here is identical to that used by an army when it threatens to shoot deserters. If all people know that these incentives hold for most others, then cooperation will not only be possible, but will be the expected norm, and the war of all against all becomes a general peace. Few contemporary political theorists think that the particular steps by which Hobbes reasons his way to this conclusion are both sound and valid. Working through these issues here, however, would carry us away from our topic into details of contractarian political philosophy. What is important in the present context is that these details, as they are in fact pursued in the contemporary debates, all involve sophisticated interpretation of the issues using the resources of modern game theory. Notice that Hobbes has not argued that tyranny is a desirable thing in itself. The structure of his argument is that the logic of strategic interaction leaves only two general political outcomes possible: Sensible agents then choose tyranny as the lesser of two evils. The distinction between acting parametrically on a passive world and acting non-parametrically on a world that tries to act in anticipation of these actions is fundamental. The values of all of these variables are independent of your plans and intentions, since the rock has no interests of its own and takes no actions to attempt to assist or thwart you. Furthermore, his probable responses should be expected to visit costs upon you, which you would be wise to consider. Finally, the relative probabilities of his responses will depend on his expectations about your probable responses to his responses. The logical issues associated with the second sort of situation kicking the person as opposed to the rock are typically much more complicated, as a simple hypothetical example will illustrate. Suppose first that you wish to cross a river that is spanned by three bridges. Assume that swimming, wading or boating across are impossible. The first bridge is known to be safe and free of obstacles; if you try to cross there, you will succeed. The second bridge lies beneath a cliff from which large rocks sometimes fall. The third is inhabited by deadly cobras. Now suppose you wish to rank-order the three bridges with respect to their preferability as crossing-points. The first bridge is obviously best, since it is safest. To rank-order the other two bridges, you require information about their relative levels of danger. Your reasoning here is strictly parametric because neither the rocks nor the cobras are trying to influence your actions, by, for example, concealing their typical patterns of behaviour because they know you are studying them. It is obvious what you should do here: Now let us complicate the situation a bit. Your decision-making situation here is slightly more complicated, but it is still strictly parametric. However, this is all you must decide, and your probability of a successful crossing is entirely up to you; the environment is not interested in your plans. However, if we now complicate the situation by adding a non-parametric element, it becomes more challenging. Suppose that you are a fugitive of some sort, and waiting on the other side of the river with a gun is your pursuer. She will catch and shoot you, let us suppose, only if she waits at the bridge you try to cross; otherwise, you will escape. As you reason through your choice of bridge, it occurs to you that she is over there trying to anticipate your reasoning. It will seem that, surely, choosing the safe bridge straight away would be a mistake, since that is just where she will expect you, and your chances of death rise to certainty. So perhaps you should risk the rocks, since these odds are much better. But wait â€¦ if you can reach this conclusion, your pursuer, who is just as rational and well-informed as you are, can anticipate that you will reach it, and will be waiting for you if you evade the rocks. So perhaps you must take your chances with the cobras; that is what she must least expect. But, then, no â€¦ if she expects that you will expect that she will least expect this, then she will most expect it. This dilemma, you realize with dread, is general: You appear to be trapped in indecision. All that might console you a bit here is that, on the other side of the river, your pursuer is trapped in exactly the same quandary, unable to decide which bridge to wait at because as soon as she imagines committing to one, she will notice that if she can find a best reason to pick a bridge, you can anticipate that same reason and then avoid her. We know from experience that, in situations such as this, people do not usually stand and dither in circles forever. However, until the s neither philosophers nor economists knew how to find it mathematically. As a result, economists were forced to treat non-parametric influences as if they were complications on parametric ones. This is likely to strike the reader as odd, since, as our example of the bridge-crossing problem was meant to show, non-parametric features are often fundamental features of decision-making problems. Classical economists, such as Adam Smith and David Ricardo, were mainly interested in the question of how agents in very large marketsâ€”whole nationsâ€”could interact so as to bring about maximum monetary wealth for themselves. Economists always recognized that this set of assumptions is purely an idealization for purposes of analysis, not a possible state of affairs anyone could try or should want to try to attain. But until the mathematics of game theory matured near the end of the s, economists had to hope that the more closely a market approximates perfect competition, the more efficient it will be. No such hope, however, can be mathematically or logically justified in general; indeed, as a strict generalization the assumption was shown to be false as far back as the s. This article is not about the foundations of economics, but it is important for understanding the origins and scope of game theory to know that perfectly competitive markets have built into them a feature that renders them susceptible to parametric analysis. Because agents face no entry costs to markets, they will open shop in any given market until competition drives all profits to zero.