Chapter 22

Nuclear Deterrence in the Algorithmic Age: Game Theory Revisited

Roy Lindelauf

Contents

22.1 Introduction.....	422
22.2 Game Theory Basics	424
22.3 Nuclear Deterrence and Basic Game Theory	427
22.4 Moving beyond the Limitations of Basic Game Theory Models	429
22.5 Nuclear Deterrence—Games and Decisions	430
References	434

Abstract Commonly used game and decision theoretic models fail to explain the empirics of deterrence. This has unjustly led many theorists to criticize the (rationality and other) assumptions underpinning of such models. No serious game theorist will contend that his theoretic model will possibly take account of all the peculiarities involved in decision making and therefore be an accurate model of such situations. Games are an aid to thinking about some of the aspects of the broader situation. Game theory models prescribe what a decision maker *ought* to do in a given situation, not what a decision maker actually *does*. To maintain nuclear strategic stability, it is of paramount importance to understand the dynamical interplay between all players involved in decision making processes with regard to nuclear strategy. History has shown some progress in understanding nuclear deterrence by the use of initial game- and decision theoretic models to alleviate the burden of human cognitive biases. Since it is highly likely that (semi-)autonomous systems will in some way participate in the future nuclear strategic landscape, combined with the fact that the nuclear deterrent decision-cycle will also be based

---

R. Lindelauf (✉)
TU Delft, Delft, The Netherlands
e-mail: rha.lindelauf.01@mindef.nl

R. Lindelauf
Netherlands Defence Academy, Breda, The Netherlands

on algorithmic analysis, rational deterrence theory is and should be an integral element of strategic thinking about nuclear deterrence. That, or it might as well be *game over*.

Keywords rationality • chicken games • tit-for-tat • autonomous weapon systems • artificial intelligence • algorithms • stability

22.1 Introduction

As discussed in the chapters by Bijlsma (Chap. 23), and Zilincik and Duyvesteyn (Chap. 24) in this volume, humans consistently and systematically make poor decisions as Nobel prize winning research on human decision making under uncertainty has shown.¹ Kahneman taught us that human thought processes as representations of the real world (so-called mental models) suffer from a multitude of cognitive biases. The mental models that humans employ are incomplete and unstable, our ability to mentally run mental models is limited and those models do not have firm boundaries and are parsimonious.² To avoid all out nuclear destruction the human calculus of deterrence should be protected from such shortcomings; deterrence theory sprang from this well. In addition, analytic techniques have been developed to alleviate the burden of cognitive biases that lead to fallacious arguments and inconsistent conclusions and should therefore be part of deterrence theoretic frameworks. A large subset of these methods entail quantitative prescriptive, descriptive and predictive models that consist of logical consistent frameworks that proceed from explicit assumptions to coherent conclusions. This chapter provides a coarse introduction into such quantitative models used to understand deterrence, with a specific focus on game- and decision theory.

During World War II quantitative modelling developed as a formal method of decision support for operations.³ The field of operations research (sometimes referred to as decision theory) and game theory that emerged from this development is rich in methods and techniques that aid higher level decision makers regarding problems concerned with the operations under their control. Simply put, operations research (OR) is the science of decision making and game theory provides an explicit normative framework on optimal decision making in either conflicting or cooperative settings. Game theory is concerned with modelling strategic interaction, hence encompasses more than one ‘player’ (whereas OR or decision theory focus on unilateral decision frameworks). For several decades such frameworks have been successfully deployed in a multitude of military strategic and operational settings. Think of the identification of resource limited interdiction actions that

---

¹See for instance Gilovich et al. 2002.

²Norman 1983.

³Washburn and Kress 2009.

maximally delay completion time of a proliferator's nuclear weapons project,⁴ dynamic task assignments for multiple unmanned combat aerial vehicles,⁵ submarine warfare,⁶ search theory⁷ and combat models,⁸ to name just a few. The application of game theory is not limited to the military but includes many sectors, be it government, business, manufacturing, healthcare, service operations, evolutionary biology, experimental sociology, psychometrics, economics or others. Game- and decision theory encompasses many different decision making situations, such as optimization of resource allocation, task allocation, coalition formation, bargaining situations, elections, signalling, pricing and of course choosing deterrent strategies.

Game theoretically speaking: deterrence equals one player threatening another player with the goal of preventing him to conduct an aggressive action that it has not yet taken (but appears willing to do). In other words, the aim of deterrence is to influence perceptions and the decision calculus of the opponent to prevent him from doing something undesired.⁹ Deterrence is therefore based on the psychological principle of a threat of retaliation. For instance, a nation wants to prevent nuclear first strikes or cyber-attacks and a company aims for the non-entry of competitors to their market. A key point in deterrence theory is *credibility*: are the threats credible or not. This depends on the attacker's beliefs on the capabilities of the defender. Clearly, any decision maker with enough concern for tomorrow is likely to be moved by deterrent threats.¹⁰ Therefore it is not surprising that deterrence is a major theme of game theory, both in economics and political science game theory plays a role in modelling deterrence. Pioneers in the application of game theory such as Thomas Schelling resorted to game theory in their discussion of nuclear deterrence even though such leading scientists were not technical game theorists per se, they simply used concepts and insights from game theory to sharpen their thinking about deterrent situations.

On the other hand, deterrence theorists trace the origin of their theories to the aftermath of World War I and classify their realist classical theory of deterrence into several strands: structural deterrence theory and decision-theoretic deterrence theory.¹¹ It is the latter theory that applies game theoretic methodology to reasoning about deterrence. In this chapter, after first introducing some basics of game theory, we will present some of the game theoretic arguments—and critiques thereof—that arise in classical deterrence theory. Next we will mention some more advanced game theoretic models that are designed to take those critiques into account. Since

---

⁴Brown et al. 2009.

⁵Duan and Yu n.d.

⁶Danskin 1969.

⁷Alpern et al. 2013.

⁸Washburn and Kress 2009.

⁹See Chap. 1 by Freedman and Chap. 2 by Mazarr in the present volume.

¹⁰Langlois and Langlois 2006.

¹¹Quackenbusch and Zagare 2016.

game theory also enters into the design and application of algorithms (of semi-autonomous systems) we will end this chapter with some observations on the recent exponential developments in computer science and the effect thereof on nuclear stability and deterrence.

In this chapter, a short introduction to normative decision making is given by presenting the basic framework of game theory. This will provide the reader with a better understanding of the standard ideas and assumptions with regard to this theory and also with some of its goals. Next several applications, including its shortcomings and advantages, of the game theory within deterrence theory are presented and discussed. The development of information technology and AI will have a large effect on nuclear security issues in the next quarter century, therefore this chapter concludes with a short outlook on future developments of nuclear deterrence with respect to algorithmic game theory in computer science in general and of artificial intelligence in particular.

22.2 Game Theory Basics

The mathematical theory of games can be divided into games of several types depending on whether,

- A. players can negotiate and form alliances or not, i.e. cooperative versus non-cooperative games,
- B. players know everything about the game (payoffs) and the other players (strategies) or not, i.e. games of (in)complete information,
- C. players act concurrently or sequentially (where each player is aware of the other player his action) or not, i.e. simultaneous versus sequential games,
- D. all the players have the same goals (are symmetric) such that only their choice of strategy determines who wins (chess) or not, i.e. symmetric versus asymmetric games,
- E. all the players have perfect information about the game (observe all the other players' moves) or not, i.e. perfect versus imperfect information games,
- F. one player's loss equals the other's gain (zero sum) or not, zero sum versus non-zero sum games.

Basically, a game involves players, strategies, payoffs and an information structure. The most well-known games are non-cooperative two player zero sum games. In general, a non-cooperative game is a sequence of moves, at each of which one of the players chooses from among several possibilities.¹² Note that some such moves may involve chance (for instance throwing a die) or are random acts of nature. At the end of the game there is some sort of payoff to all of the players. This for instance can be money, satisfaction, or any other quantifiable variable. In

---

¹²Owen 2001.

general, non-cooperative games are modelled either in *extensive*- or *normal* form. The former includes the possibility of alternation of moves by players and situations where players can have less than perfect information such as not knowing other player's payoffs or possible moves. The latter involves the assumption that, given knowledge of the game and its payoffs, each player has already decided what he will do before the game starts, i.e. each player chooses a strategy before the game and they do so simultaneously. This may be a restrictive assumption at first, but it encapsulates the idea of devising a plan ('strategy' in game theoretic nomenclature) for a coming situation.

Most often the game theorist is interested to devise the best possible plan for a given game, i.e. to find optimal strategies for each player. Optimality consists of maximizing the payoff to the respective players and looking for equilibrium situations. Simply put, an equilibrium occurs if each player is satisfied. Below simple examples of an extensive form game and a normal form game are given. The main difference between a game in extensive form or normal form is in the sequentially to a player's moves. The former allows for players to move after each other such that players can observe the other's moves, the latter assumes that players decide upon optimal strategies before the game commences.

In Fig. 22.1 an example of an extensive-form game with two players (player 1 and player 2) is shown. Many examples in international relations theory are modelled by simple extensive-form games.¹³ The game in Fig. 22.1 commences at the root of the tree where player 1 can choose between option Z or W. Depending on player 1's moves, player 2 can either choose between A and B or between C and D. In the former case player 1 again is presented with two options: X or Y. The dotted line indicates an *information set*, i.e. it exemplifies a situation of incomplete information: player 1 (noted above the dotted line) cannot distinguish between the states in the information set, that is, he does not know whether player 2 will choose A or B. Finally, the numbers denoted at the terminal vertices indicate the payoff to the respective players.

To infer optimal options for the players in this game, to 'solve' an extensive-form game, several solutions concepts exist in game theory. The most well-known solution to an extensive game is called backward induction where one reasons backward in time to solve each subgame, reasoning from the optimality of the previous solved subgame.¹⁴ It is a theorem in game theory that a subgame perfect Nash equilibrium can always be obtained by backward induction in finite perfect information games. Even though the game in Fig. 22.1 is not of the perfect information type, it still can be solved using this procedure (due to its payoff structure). In the game in Fig. 22.1, starting at player 1 in his choice between X or Y it can be seen that X always favors Y ($5 > 3$ and $4 > 2$) even though player 1 does not know player 2's choice between A and B. Player 2 then (knowing that player 1 will choose X) favors B over A ($6 > 0$), additionally player 2 prefers C over D ($4 >$

¹³See Langlois and Langlois 2006 regarding the Cuban missile crisis.

¹⁴Neumann and Morgenstern 1944.

Fig. 22.1 An example of an extensive-form game (Source Roy Lindelauf)

2). Finally, player 1's choice at the root is between Z (which will yield 6) or W (which will yield 4). Hence the solution obtained by backward induction yields player 1 choosing Z, X and player 2 choosing C, B. Clearly more realistic games contain more players, more moves and more options per player. If the game has a finite horizon and is of perfect information the solution procedure as sketched above remains the same (and can be computed algorithmically).

Another very common approach to model strategic interaction with game theory is according to the *normal-form*. Informally a normal-form game consists of players each of which have strategies that after selection are played *simultaneously*. Each strategy selection for all players results in payoffs (to each player) that can be observed by all players. Those payoffs are denoted in a *payoff matrix* (see Fig. 22.2) which can be analysed using well defined concepts (such as dominating strategies, pure or mixed Nash equilibria, etc.). Solutions of a normal-form game then consist of *good* strategy prescriptions for each player. When the game situation repeats over time and/or space those solutions come in the form of probability distributions over the set of pure strategy options for each player. The most well-known such solution is the Nash equilibrium which consists of the strategy profile for all players such that no player can unilaterally benefit from deviation from that profile. If the game is of complete information (each player knows all the options of every player and all corresponding payoffs) all players can compute the optimal strategy of each player. However, even though each player thus knows the optimal play of the other players, he/she still does not know the actual option those players will play (because those strategies are given probabilistically). Hence normal-form games provide optimal plays that are rational but unpredictable.

Consider the following simplified introductory example. Somewhere on a remote island drug smugglers regularly drop off small shipments of illegal drugs at either one of two locations, A or B. The police unit on the island has very limited resources and can observe only one location at a time. Knowing that the drop-off capacity at location A is twice that of B, the question arises which location should be observed more often (and how often). Similarly, the smugglers wonder at which location to drop their drugs and with what frequency.

Fig. 22.2 An example of a normal-form game (*Source* Roy Lindelauf)

		Smugglers	
		A	B
Police	A	2,-2	0,0
	B	0,0	1,-1

The Nash equilibrium to this game equals the Police observing location A with probability 1/3 and location B with probability 2/3 (due to the symmetry of the game the same holds for the smugglers). In practice this translates to the police throwing a regular dice before each observation and when it lands on either 1, 2, 3 or 4 the observation takes place at location B, otherwise at location A (similar for the smugglers). The expected average capacity (kilos of cocaine) of seized drug shipments then equals 2/3, i.e. if the police did 100 observations (and each shipment contained either 1 or 2 k of cocaine) according to this strategy then on the average there would be a total of about 66 k of cocaine seized. Clearly normal-form games that model reality more realistically contain more than two options or two players. The normal-form game theoretic framework however has been very successful and applied in a plethora of applications.

22.3 Nuclear Deterrence and Basic Game Theory

Initial game theoretic models of deterrence are extreme simplifications of the complicated reality of deterrent situations. In general, this contention holds for many normal and extensive form games used in international relations theory. Often such text book models are two person games with both players only having two options (the normal-form as presented in the previous paragraph). Perhaps the most well-known game in IR theory is the prisoner’s dilemma, used to model the Cuba crisis for instance, developed in the 1950s by RAND researchers.¹⁵ Such two-by-two games serve as gentle introductions into the ideas and concepts of basic game theory, but lack depth and structure for modelling realistic situations.

An example of a two-by-two game to model nuclear crisis is the *chicken game*.¹⁶ This game represents the situation where two teenagers are speeding towards each other in their cars at the middle of the road, i.e. representing two nuclear belligerents threatening each other with all-out nuclear war. Each player only has two options: to swerve (S; not attack) or to non-swerve (NS; attack). The corresponding payoff structure is modelled as follows: the first player to swerve loses. However, this loss is not as bad as both players not swerving (resulting in mutual destruction).

¹⁵Dresher 1961.

¹⁶Kahn 1960.

Clearly, a player prefers the situation of both players swerving above the situation where he swerves and the other does not (better to have no nuclear strikes than to be destroyed by one). This results in the following ordinal preference structure (for player 1): $(NS, S) > (S, S) > (S, NS) > (NS, NS)$. Because the game is completely symmetric the same holds for the second player and it can easily be seen that only if both players cooperate a compromise could emerge. Otherwise, if only one player cooperates (agrees not to strike) the other player can exploit this (by striking). This game theoretic equilibrium outcome does not represent reality.¹⁷

It is easy to argue that the chicken game abstracts away too much aspects of a nuclear crisis. The game assumes that both players only have two options, that they determine their strategy beforehand, that there is no observation of others' actions and that each player has complete information about the game. Additionally, empirical evidence shows that, in the case of approximately equal opponents, it is better for a player to escalate when challenged with a nuclear strike than to submit.¹⁸ Chicken games do not allow for such step-by-step iterations.

In Quackenbusch and Zagare (2016) a simple extensive-form game to model deterrence is introduced, the rudimentary asymmetric deterrence game (RADG). This game consists of two players (challenger and defender) and both players have two options at their disposal: cooperate and defect for challenger and concede and deny for defender. Challenger moves first. It is stated that the assumption that conflict is the worst outcome is 'the defining assumption of decision-theoretic decision theory', translated into the lowest payoff for both players in case of 'conflict'. Then it is reasoned that, by use of RADG, this leads to the paradox of deterrence: the contention that bilateral relationships between nuclear equals are stable even though the 'solution' of the RADG does not equal the status quo. The solution of the RADG as presented by Quackenbush and Zagare, i.e. the Nash equilibrium obtained by backward induction, equals the situation where challenger chooses 'defect' and defender chooses 'concede'. Indeed, this is not a stable bilateral situation between nuclear equals. This paradox is easily resolved by changing the payoff structure of the game, hence it is not a shortcoming of decision theory but rather of modelling choice.

Summarizing, early game theoretic models used by modellers of deterrence lack complexity to include,

1. situations of escalation, i.e. players react to each other inducing continuous iterations of developing situations,
2. attackers and defenders (almost always) are not exactly aware of each other's strategy options and utility calculations (and there can be more than two players), i.e. incompleteness of information dominates international political decision making regarding deterrence,

¹⁷Another (mixed) strategy equilibrium exists in chicken games but that is not mentioned here because it does not add to the discussion.

¹⁸Zagare 1987.

3. attackers and defenders are not exactly aware of the moves other players have made (and there can be more than two options per player).

Research in game theory—outside of the scope of deterrence—recognized all of the restrictions mentioned above.

22.4 Moving beyond the Limitations of Basic Game Theory Models

A plethora of advancements have been made to overcome those limitations. The easiest: having more than two options for each player and having more than two players. Additionally, iterated games (also called ‘repeated games’) were developed to analyse series of decisions that are not ‘one-shot’; they overcome the first objection as mentioned above. Knowing that a game will continue indefinitely will impact how players choose their strategies because players have knowledge of the past behaviour of their rivals (they observe their choices). The Soviet-US arms race for instance has been most commonly modelled as an iterated prisoners’ dilemma [IPD] (Majeski 1984). This led to the famous TIT-FOR-TAT (TFT) decision rule that consists of choosing ‘cooperate’ during the first iteration and then copying what the other player did in the previous round thereby rewarding cooperative behaviour and punishing otherwise. It turned out that this decision rule did surprisingly well in many comparisons of strategies for the IPD because of its properties of niceness, forgiveness and retaliatoriness. This results in a model where on any given trial both superpowers are better off arming regardless of what the other side chooses, but if both sides arm the outcome is less desirable than had both sides reduced their supply of weapons.¹⁹

Several extensive form models have been introduced, such as Hawks and Doves games that include incomplete information situations and elements of nuclear brinkmanship by introducing escalation ladder models²⁰. However, this model still suffers from many of the shortcomings mentioned above. It was John Harsanyi who first developed game theoretic models to deal with situations of incomplete information.²¹ With respect to deterrence for instance the attacker’s beliefs on the credibility concerning the defender’s deterrent threat are uncertain. Such incomplete information could be about the other player’s motivations, strategy options, resolve, beliefs about the other player and others aspects.

Clearly the problem of deterrence has also inspired more advanced forms of game theory. Nobel Prize winner Robert Aumann together with Michael

---

¹⁹Plous 1993.

²⁰See for instance Langlois’ online chapter three of *Applicable Game Theory*.

²¹Harsanyi 1967.

Maschler²² wrote a book on the application of mathematical utility theory to disarmament. They formulated repeated two-player games in which one (or both) of the players lack complete information on the payoffs in the stage-game matrix. They showed that when one of the two players has special information not available to the other, then he can use this information to his advantage only to the extent that he reveals it. Using several theorems, they showed that optimal strategies in repeated games of incomplete information contain certain interesting peculiarities which are best illustrated by the following analogy: consider player 1, a policy maker who does not play the game himself, he uses a negotiator to play the game for him instead. Aumann and Maschler showed that the optimal strategy for the policy maker is to fool his negotiator to the extent that he reveals him a certain amount of information on how to negotiate (the type of negotiator he is) according to some probability distribution (determined by mathematical analysis). The interesting fact is that complete disclosure nor complete concealment of secret information from one's negotiator is in general an optimal strategy and that there exists a random mechanism that describes exactly what partial information should be disclosed to the negotiator.²³

22.5 Nuclear Deterrence—Games and Decisions

The preceding paragraphs illustrated that commonly used game and decision theoretic models fail to explain the empirics of deterrence. This has unjustly led many theorists to criticize the (rationality and other) assumptions underpinning of such models.²⁴ Next to the reasons already mentioned, no serious game theorist will contend that his theoretic model will possibly take account of all the peculiarities involved in decision making and therefore be an accurate model of such situations. Games are an aid to thinking about some of the aspects of the broader situation. The corresponding conclusions therefore will reflect general insights that can be useful in the weighing of multiple criteria upon making a decision. Game theory models prescribe what a decision maker *ought* to do in a given situation, not what a decision maker actually *does*.

Much in the same way, decision theory for instance has taught us by mathematical analysis that commonly accepted beliefs about decision procedures with three or more candidates will *always* lead to a dictator; by listing basic properties of decision methods satisfied by all (democratic) election methods Arrow showed this

---

²²Aumann and Maschler 1995.

²³For a much more in-depth analysis of these assertions, we refer to Aumann and Maschler 1995.

²⁴Lebow and Stein 1989.

in his famous theorem.²⁵ This shows the fallacy of human reasoning and the necessity of logical consistent thinking; informal arguments can lead to seemingly correct conclusions which in reality contain falsehoods. Such contentions in all likelihood also hold for arguments in deterrence theory. The theories of games and decisions are therefore of innumerable value to provide a coherent explicit framework and to alleviate the burden of cognitive biases in decision settings such as deterrence. No framework, be it quantitatively motivated or not, will ever explain all the peculiarities encompassed in complex deterrence settings. ‘All models are wrong, but some are useful’ as famous statistician George Box used to say.²⁶ So what then is the future of game theory in deterrence?

First, game theory can help to lay an axiomatic foundation under the theory of deterrence, much as decision theory did for the theory of democratic elections (see our earlier mention of Arrow’s theorem). Second, the world is witnessing unprecedented technological innovations in information technology. Algorithms are entering each and every aspect of our lives, from choosing which movie to watch at night to predicting poaching of wildlife. The exponential growth of processor speed, data storage, computational analysis and technology in general are changing the future battlefield. Systems embedded with algorithms that make decisions on its behaviour are commonplace and are expected to proliferate in the future. It comes as no surprise that these advancements in computer science enable *rational* decision making within the field of deterrence along another avenue of approach. The future battlefield will see a mix of (semi-)autonomous weapon systems with manned systems.²⁷ It is highly likely that these systems will deploy game- and decision theory based algorithms to coordinate and control.²⁸ Autonomous weapon systems base their decisions on all kinds of algorithms. These artificial intelligence and autonomous systems have the potential to dramatically affect nuclear deterrence and escalation.²⁹ The speed of decision making, its differences from human understanding, the willingness of many countries to use autonomous systems, our relative inexperience with them, and continued developments in these capabilities are among the reasons.³⁰ A similar situation has already been witnessed in the field of stock trading where high frequency automatic trading algorithms are deployed to conduct autonomous trading. This contributed to the flash crash of the stock market in 2010 where computers in fast automated markets made buy-sell decisions in fractions of seconds.³¹

Game and decision theoretic concepts often translate directly into such algorithms. Actually game- and decision theory is an integral element of artificial

---

²⁵Saari 2001.

²⁶Box 1976.

²⁷Wong et al. 2020.

²⁸Marden and Shamma 2018; Morgan et al. 2018.

²⁹See Chap. 21 by Wilner and Casey in the present volume.

³⁰Wong et al. 2020.

³¹Kirilenko and Samadi 2017.

intelligence.³² Machine learning classifiers such as a support vector machine for instance can be seen as strategic two player games, i.e. one player is challenging the other in finding the optimal hyperplane by giving him the most difficult points to classify. Many algorithms implemented by (semi-)autonomous systems are based on rational decision making. In multi-agent reinforcement learning for instance agents *learn* by interacting with the environment and with other agents. Often the Nash equilibrium represents the collaboration point between the different agents (players). In short, this forces game theory in the future of nuclear deterrence along several avenues of approach. Below we exemplify three of them.

- A. The design of nuclear weapon decision support algorithms, for instance with respect to the detection and tracking of adversary launchers for counterforce targeting;
- B. With respect to coordination and competition between (semi-)autonomous nuclear systems, for example consider Russia's nuclear powered undersea drone that can carry a thermonuclear warhead and that should be able to operate autonomously for prolonged periods of time;³³
- C. Regarding (adversarial attacks on) algorithms used in the nuclear infrastructure, for instance by data poisoning corresponding SCADA systems.³⁴

First, consider one of many nuclear weapon decision processes: the targeting process. This is the practice that aims at achieving specified effects on and beyond the battlefield that employs classic kinetic lethal actions as well as non-military, non-kinetic, and nonlethal activities. The process consists of six phases of which the second phase—*target analysis, vetting, validation, nomination and prioritization*—is clearly of interest to automation of (nuclear weapon) decision support. With the massive increase of data in the Intelligence Surveillance and Reconnaissance (ISR) domain comes the need of automated analysis simply because the amount of data exceeds the timely analysis capacity of human analysts. The second phase of the targeting process can benefit from the use of automated analysis since it provides opportunities to deal with the complexity, scope and scale of the targeting process.³⁵ Decision support algorithms for nuclear weapon targeting come in many shapes and forms and can benefit from game theoretic approaches. Target prioritization for instance consists of ranking targets because resources are scarce. This is related to solution concepts in cooperative game theory such as the Shapley value that axiomatically defines a formula to derive the power of 'players' that create value upon cooperation. With respect to nuclear targeting this relates to the importance of a target with respect to the value of a subset of targets that can be engaged given cost and capacity restrictions. Power indices in cooperative game

---

³²Norvig and Russell 2016.

³³Geist and Lohan 2018.

³⁴Terziyan et al. 2018.

³⁵Ekelhof 2018.

theory provide a sound basis to support such decisions and are applied in a plethora of security domains.³⁶

Second, consider coordination and competition between (semi-)autonomous systems, i.e. the field of multi-agent systems (MAS)—an area in distributed artificial intelligence—that consists of multiple autonomous interacting units each with their own sensory systems and goals. Based on resources and agents skills MAS systems will either be in cooperation and collaboration or competition.³⁷ Military applications of MAS frameworks for instance consist of surveillance, navigation and target tracking and are clearly also beneficial in nuclear settings. Future systems like the Russian undersea drone for instance have to operate autonomously to achieve individual goals over long periods of time and are expected to interact with other agents that influence each other's decisions. One advantage of such a drone system is its capability for ultra-long loitering periods as there is no human crew that needs time to recuperate and recover. Therefore, it also needs to be equipped with smart decision procedures. One possible approach to develop such protocols is by multi-agents reinforcement learning, a research area within AI that uses game theory to learn optimal behaviour of agents through trial and error interaction with the environment and with other agents. In such a setting agents are assumed to be players in a normal-form game which is played repeatedly.³⁸ The importance of understanding the dynamics of such game theoretic algorithms is evident and still an active field of open research.

Third, future AI developments might put the nuclear infrastructure even more at risk in various ways. Inadvertent nuclear escalation is being driven by the fact that nuclear command, control, communications, computers, intelligence, surveillance and reconnaissance (C4ISR) capabilities are entangled with nonnuclear weapons.³⁹ Cyber operations, empowered by AI algorithms, magnify and aggravate the challenges associated with C4ISR as military cyber offensives threaten the elimination of C4ISR capabilities.⁴⁰ Hence dual-use C4ISR capability could become under attack during a conventional conflict and prove escalatory in the nuclear domain. Protecting such critical infrastructure can be done by assisting decision systems using game theoretic models that compute optimal defender strategies in near real-time, thus providing efficient ways of allocating scarce resources for defence. An example of such a model for instance consists of power grid defence against malicious cascading failures.⁴¹ Another area where game theory meets artificial intelligence is the field of generative adversarial networks. Here deep learning tasks

---

³⁶Van Campen et al. 2018.

³⁷Parker 2008.

³⁸Nowé et al 2012.

³⁹Acton 2018.

⁴⁰Futter 2016; Li 2018; Sweijs and Osinga 2020.

⁴¹Shakarian and Lindelauf 2014.

can be viewed as strategic games.⁴² Such models have also been used in data poisoning attacks that target machine learning algorithms by injecting malicious data-points into the training dataset.⁴³ Modern communication technologies used in SCADA systems that operate nuclear infrastructure introduce security vulnerabilities such as data poisoning of their AI driven decision support algorithms.⁴⁴

To maintain nuclear strategic stability, it is of paramount importance to understand the dynamical interplay between all players involved in decision making processes with regard to nuclear strategy.⁴⁵ History has shown some progress in understanding nuclear deterrence by the use of initial game- and decision theoretic models to alleviate the burden of human cognitive biases. Since it is highly likely that (semi-)autonomous systems will in some way participate in the future nuclear strategic landscape,⁴⁶ combined with the fact that the nuclear deterrent decision-cycle will also be based on algorithmic analysis, rational deterrence theory is and should be an integral element of strategic thinking about nuclear deterrence. That, or it might as well be *game over*.

References

- Acton J M (2018) Escalation through Entanglement. *International Security*, 43.1:56–99
- Alpern S, Fokkink R, Gasieniec L, Lindelauf R, Subrahmanian VS (2013) *Search Theory: A Game Theoretic Perspective*. Springer, New York
- Aumann R J, Maschler M B (1995) *Repeated Games with Incomplete Information*. MIT Press, Massachusetts Institute of Technology
- Box GEP (1976) Science and Statistics. *Journal of the American Statistical Association* 71.356:791–799
- Brown G et al (2009) Interdicting a Nuclear-Weapons Project. *Operations Research* 57.4
- Danskin JM (1969) A Helicopter versus Submarine Search Game. *Operations Research*, Vol. 16, No.3
- Dresher M (1961) *The Mathematics of Games of Strategy: Theory and Applications*. Prentice-Hall, Englewood Cliffs NJ
- Ekelhof M (2018) Lifting the fog of targeting. *Naval War College Review*, 71.3. US Naval War College Press
- Futter A (2016) *Cyber Threats and Nuclear Weapons. New Questions for Command and Control, Security and Strategy*. RUSI Occasional Paper
- Geist E, Lohan A J (2018) *How Might Artificial Intelligence affect the Risk of Nuclear War?* RAND Corporation, Santa Monica CA
- Gilovich T, Griffin D, Kahneman D (2002) *Heuristics and biases: The psychology of intuitive judgment*. Cambridge University Press, New York

---

⁴²Tambine 2019.

⁴³Munoz-Gonzalez et al. 2019.

⁴⁴Munoz-Gonzalez et al. 2019.

⁴⁵Morgan et al 2017.

⁴⁶Horowitz 2019.

- Gomez Rivera AO et al (2019) Towards Security and Privacy of SCADA Systems through Decentralized Architecture. International Conference on Computational Science and Computational Intelligence (CSCI), IEEE Xplore
- Haibin D et al (2015) A Predator Prey Particle swarm optimisation approach to multiple UCAV air combat modeled by dynamic game theory. IEEE/CAA 2.1
- Harsanyi JC (1967) Games with incomplete information played by “Bayesian players”, I-III. part I. The Basic Model. *Management Science* 14.3:159–182
- Horowitz M C (2019) When speed kills: Lethal autonomous weapons systems, deterrence and stability. *Journal of Strategic Studies* 42.6
- Kahn H (1960) On thermonuclear war. Routledge, New York
- Kahneman D, Tversky A (1974) Judgement and Uncertainty: Heuristics and Biases. *Science*, 27.185
- Kirilenko A S, Samadi M (2017) The flash crash: High-frequency trading in an electronic market. *The Journal of Finance*, Wiley Online Library
- Kopelman S (2020) Tit for Tat and Beyond. *The Legendary Work of Anatol Rapoport. Negotiation and Conflict Management Research* 13.1:60–84
- Langlois J P, Langlois C (2006) Bargaining and the Failure of Asymmetric Deterrence: Trading off the Risk of War for the Promise of a Better Deal” (with Catherine Langlois), *Conflict Management and Peace Science* 23:159–180
- Lebow N R, Stein J G (1989) Rational Deterrence Theory: I Think, Therefore I Deter. *World Politics* 41.2:208–224
- Li J (2018) Cyber security meets artificial intelligence: a survey. *Frontiers of Information Technology & Electronic Engineering* 19:1462–1474
- Majeski SJ (1984) Arms races as iterated prisoner’s dilemma games. *Mathematical Social Sciences* 7(3):253–266
- Marden JR, Shamma JS (2018) Game Theory and Control. *Annual Review of Control, Robotics, and Autonomous Systems* 1:105–134
- Morgan FE et al (2018) Military Applications of Artificial Intelligence. *Ethical Concerns in an Uncertain World*. RAND Corporation, Santa Monica CA
- Morgan GP et al (2017) Sociocultural Models of Nuclear Deterrence. *IEEE Transactions on Computational Social Systems* 4.3
- Munoz-Gonzalez L et al (2019) Poisoning Attacks with Generative Adversarial Nets. [arXiv:1906.07773v2](https://arxiv.org/abs/1906.07773v2)
- Neumann J, Morgenstern O (1944) *Theory of Games and Economic Behavior*. Princeton University Press, Princeton
- Norman DA (1983) Some Observations on Mental Models. In: Gentner D, Stevens A (eds) *Mental Models*. Lawrence Erlbaum Associates Cognitive Science Series, Mahwah NJ
- Norvig P, Russell S (2016) *Artificial Intelligence: A Modern Approach*. Pearson Education Limited
- Nowé A, Vrancx P, De Hauwere YM (2012) Game Theory and Multi-agent Reinforcement Learning. In: Wiering M, van Otterlo M (eds) *Reinforcement Learning. Adaptation, Learning, and Optimization*, vol 12. Springer, Berlin/Heidelberg
- Owen G (2001) *Game Theory*. Academic Press
- Parker L E (2008) Distributed Intelligence: Overview of the field and its application in multi-robot systems. *J. Phys. Agents* 2.1:5–14
- Plous S (1993) The nuclear arms race: Prisoner’s dilemma or perceptual dilemma. *Journal of Peace Research*
- Quackenbusch S L, Zagare F C (2016) Modern Deterrence Theory: Research Trends, Policy Debates, and Methodological Controversies. *Oxford Handbooks Online*
- Saari DG (2001) *Decisions and Elections; Explaining the Unexpected*. Cambridge University Press, Cambridge
- Shakarian S, Lindelauf R (2014) Power Grid Defense Against Malicious Cascading Failure. 13th Conference of Autonomous and Multiagent Systems, AAMAS-14
- Sweijts T, Osinga F (2020) Maintaining Nato’s Technological Edge. *Whitehall Papers* 95.1

- Tamine H (2019) Deep Learning Meets Game Theory: Bregman-Based Algorithms for Interactive Deep Generative Adversarial Networks. *IEEE Transactions on Cybernetics* 50.3
- Terziyan T et al (2018) Industry 4.0 Intelligence under Attack: From Cognitive Hack to Data Poisoning. In: Dimitrov K (ed) *Cyber Defence in Industry 4.0 Systems and Related Logistics and IT Infrastructures*. IOS Press, Amsterdam
- Van Campen T, Hamers H, Husslage B, Lindelauf RA (2018) New approximation method for the Shapley value applied to the WTC 9/11 attack. *Social Network Analysis and Mining* 8.3
- Washburn A, Kress M (2009) *Combat Modeling*. Springer, New York
- Wong YH et al (2020) *Deterrence in the Age of Thinking Machines*. RAND Corporation, Santa Monica CA
- Zagare F (1987) *The Dynamics of Deterrence*. University of Chicago Press, Chicago

Roy Lindelauf is assistant professor of Quantitative Intelligence Analysis at the Department of War Studies of the Netherlands Defence Academy. A former Apache Attack Helicopter pilot, his research includes Quantitative Intelligence Analysis, Game Theory, Network Analysis, Data Mining and Algorithms. He obtained his PhD from Tilburg University.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.