DECEPTION IN GAME THEORY: A SURVEY AND MULTIOBJECTIVE MODEL

THESIS

Austin L Davis, Capt, USAF

AFIT-ENG-MS-16-M-011

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense or the United States Government. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

AFIT-ENG-MS-16-M-011

DECEPTION IN GAME THEORY
A SURVEY AND MULTIOBJECTIVE MODEL

THESIS

Presented to the Faculty
Department of Electrical and Computer Engineering
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
in Partial Fulfillment of the Requirements for the
Degree of Master of Science in Computer Science

Austin L Davis, B.S. Mathematics
Capt, USAF

March 2016

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

AFIT-ENG-MS-16-M-011

DECEPTION IN GAME THEORY
A SURVEY AND MULTIOBJECTIVE MODEL

THESIS

Austin L Davis, B.S. Mathematics
Capt, USAF

Committee Membership:

Dr. B. J. Borghetti
Chair

Dr. G. B. Lamont
Member

Maj B. R. Woolley, PhD
Member

AFIT-ENG-MS-16-M-011

Abstract

Game theory is the study of mathematical models of conflict. It provides tools for analyzing dynamic interactions between multiple agents and (in some cases) across multiple interactions. This thesis consists of two scholarly articles that address deception from a game theoretic (GT) perspective.

The first article is a survey of GT models of deception. The survey describes the ways researchers use game theory to measure the practicality of deception, model the mechanisms for performing deception, analyze the outcomes of deception, and respond to, or mitigate the effects of deception. The survey highlights several gaps in the literature. One important gap concerns the benefit-cost-risk trade-off made during deception planning.

To address this research gap, the second article introduces a novel approach for modeling these trade-offs. The approach uses a GT model of deception to define a new multiobjective optimization problem called the deception design problem (DDP). Solutions to the DDP provide courses of deceptive action that are efficient in terms of their benefit, cost, and risk to the deceiver.

A case study based on the output of an air-to-air combat simulator demonstrates the DDP in a 7×7 normal form game. Two prominent features are observed in the solutions. First, many of the resulting solutions are zero-risk. A zero-risk solution implies one of two properties: either 1.) the deceived player has no recourse if the deception is discovered, or 2.) the deception cannot be revealed unless the deceived player intentionally adopts a strategy that he perceives to be suboptimal. Second, the solutions tended to distort a considerable portion of the game's payouts: at least 94 of the 98 payouts were modified in each of the seven efficient solutions. Several techniques are described that may reduce the number of changed payouts.

This thesis makes several significant contributions to the literature:

- A survey of the GT models of deception
- Introduces a GT model for environmental deception in normal form games
- Defines a quantitative measure of deceptive risk when employing environmental deception
- Introduces the first multiobjective GT model of deception
- Introduces the first GT model that addresses the benefit, cost, and risk trade-off inherent to every deception
- Marks the first time a multiobjective evolutionary algorithm has been applied to solve a GT problem

Contents

Page

Page

List of Figures

List of Tables

DECEPTION IN GAME THEORY

A SURVEY AND MULTIOBJECTIVE MODEL

I. Introduction

1.1 Overture

Numerous historical conflicts relied on deception to achieve victory. The success of the Trojan horse in the capture of Troy is a well-known example. In the past, individual commanders planned and executed tactical-level deceptions, but by the early 20th century, warfare had become prohibitively complex. The Prussian general Clausewitz (1780-1831) argued that even in the absence of deception it is hard to see through the fog of war; without an accurate view of the battlefield, it is impossible to anticipate the adversary's response to an elaborate deception. "To prepare a sham action with sufficient thoroughness to impress an enemy requires a considerable expenditure of time and effort.... there is always the risk that nothing will be gained and that the troops deployed will not be available when they are needed" [81].

In response to warfare's growing complexity, deception planning transitioned from individual commanders to strategic planning departments. These departments developed unified deceptive plans across multiple battlefields, e.g. Operation Bodyguard ahead of the Allied landing at Normandy. However, modern warfare is characterized by both an increase in complexity as well as an accelerated operations tempo.

The most extreme examples unfold in the domain of cyberspace, where attackers can compromise a network without ever alerting a network administrator. Concerning surprise in the traditional warfighting domains, Clausewitz argued, "the enemy forces can never assemble and advance so secretly that the defender's first news of it would come from his outposts" [81]. Yet, this is exactly the situation faced by cyberspace operators. Historically, cyber network defense relied heavily on human intervention to detect malicious activity in a network [35]; however, even against a well-defended network, attackers can exploit unforeseen vulnerabilities that human administrators overlook. It is estimated that 35 percent of all cyber attacks are never detected [62].

Recent, high-profile cyber-attacks [65] have highlighted a critical need for defenses against sophisticated, well-organized adversaries. Hamilton, et al. [35] envisioned the application of game theoretic techniques to combat this growing threat, noting that game theory provides a means for evaluating hundreds of thousands of possible scenarios and recommending several responses to adversarial activities. Furthermore, since cyberspace exists as a virtual world, agents have much greater influence over their adversary's perception of the environment. Cyberspace, therefore, offers tremendous opportunities for deception planning and execution by autonomous agents. In general, any time the decision-to-action-to-effect chain is so small that it is difficult for a human operator to react, deception can be an effective tool. Thus, although this research is conceived with cyberspace in mind, the contributions apply to a broad range of situations faced by the United States military.

1.2 Contribution

This research is presented as a pair of scholarly articles. The first article is located in Chapter II, and the second article is contained in Chapter III. The first article surveys the corpus of game theoretic models of deception. The main contributions of the survey are:

- To develop a case for the applicability of game theory to deception planning and modeling

- To describe the ways in which game theory is used to measure the practicality of using deception
- To describe the mechanisms for performing deception (e.g. camouflage) in game theoretic terms
- To survey the literature on game theoretic models of deception and summarize the conclusions
- To identify gaps in the body of research that require additional attention

The most significant gap identified in the survey is the frequent omission of cost and risk measures. Few studies addressed cost as an explicit consideration when planning a deception. No study addressed the risk of being countered by a deception-aware opponent.

To address this gap, the second article presents a novel approach that combines multiobjective optimization, game theory, and risk assessment principles to the problem of planning efficient deceptions. The intent is to evaluate the desirability of inducing a change in the opponent's perception of the conflict. An objective function is defined for each goal—benefit, cost, and risk—and instances of the Deception Design Problem are solved using a multiobjective evolutionary algorithm (MOEA). This article has three major contributions. First, it is the first effort which quantitatively measures benefit, cost, and risk in a single game theoretic model. Second, it is the first effort that uses MOEAs to solve a game theoretic problem; however, the converse is common. (Game theory has been used to solve multiobjective problems.) Third, the approach can be adapted and used by practitioners to inform deception design processes.

II. Modeling Deception using Game Theory and its Extensions: A Survey

2.1 Introduction

Deception is a deliberate activity executed to mislead others into taking specific actions that benefit the deceiver. Decision making in deceptive environments is similar to decision making under uncertainty. However, deceptions introduce additional, unique challenges because they are engineered and adversarial. One technique to overcome uncertainty is to gather additional information about a conflict. However, in deceptive conflicts, this extra information can be manipulated to cause increased damage to the target decision maker (called the *mark*). The key difference is that deceptions are not the product of random chance or natural phenomenon; they are designed. This difference distinguishes deceptive situations from non-deceptive situations.

With the tremendous potential for gain and loss, many have undertaken the study of deception. Researchers have studied deception in fields as diverse as philosophy, law, psychology, and sociology. In economics, deception is used as a strategic tool to improve corporate performance [49]. Although deception in social and economic interactions can be destructive, it has long been employed as a tool in warfare. Sun Tzu stated that all warfare is based on deception, and its usage has been described in battles from the conquest of Canaan to World War II [17, 77] and its employment continues into the emerging domain of cyberspace [74].

Deception has been observed in the animal kingdom, as well. Many animals use mimicry to ward off predators: the red milk snake mimics the color pattern of the highly venomous coral snake. It is from this mimicry that the saying is derived, "Red on black, venom lack..." Some animals use complex motion to deceive their prey and potential mates. Dragonflies, for instance, can plan their flight path to appear to their target as if they are motionless [55]. The only evidence of their approach is their looming size. By using this *motion camouflage*, the dragonfly can approach its target unnoticed.

The fact that deception is not simply a human phenomenon indicates that there may be generalizable principles that govern its operation [39]. It is, therefore, no surprise that researchers apply mathematics to the study of deception. Game theory provides many tools to analyze mathematical models of conflict including those where uncertainty and deception are used. Much work has been done to study deception using game theory, and it is appropriate to summarize the themes and characteristics observed in the literature.

This chapter is a survey of the game theoretic study of deception. The remainder of this chapter is organized as follows: Section 2.2 provides an overview of game theory. Section 2.3 describes the necessary conditions for deception, its desirable properties, and summarizes the various deceptive mechanisms observed in the literature. Several new deceptive mechanisms are recommended, as well. Section 2.4 highlights a need for additional work focused on detecting deception in games. Section 2.5 surveys the various deception models. Section 2.6 describes deception mitigation and counter-deception in the literature. Concluding remarks and future work are presented in Section 2.7.

2.2 Background

Game theory is a field of mathematics that models conflict as a game. In a game, the players compete simultaneously to attain the best outcome. In 1928, John von Neumann authored his seminal paper, *On the theory of games*, and set apart game theory as a unique field of study [82]. Later, he provided an axiomatic definition of the theory in his book co-authored with Oskar Morgenstein, *Theory of Games and Economic Behavior* [83]. Myerson [56] defines game theory as "the study of mathematical models of conflict and cooperation between intelligent, rational decision-makers." A game captures some of the complexity in a model—a model that can be used to identify desirable strategies for the real conflict. Often, there is no single *best* strategy. Instead, a game theorist recommends a mixed-strategy where players select their actions according to a probability distribution. Modern game theory began by considering the existence of mixed-strategy equilibria in two-player zero-sum games, i.e. in games where the losses of one player are the gains of another player. In its simplest form, a game consists of a collection of players, the set of actions available to each player, and the outcomes that result from their combined actions. Each outcome is assigned a vector, called the *payouts*. When the combined actions of the players result in an outcome, each player is rewarded according to their index in the payout vector, i.e. player 1 receives the value in positon 1, player 2 receives the value in position 2, and so on. These elements are used to determine how each player should play the game, i.e. the players' strategies. A Nash equilibrium is a solution concept in which no player has an incentive to deviate unilaterally from a particular strategy. The Nash equilibrium is named in honor of John Nash who in 1951 proved that every non-cooperative game with finite action set exhibits at least one such equilibrium [57].

Models of Information.

Aumann and Maschler [3, p. 66] emphasize two ways in which real-life conflicts differ from traditional game theoretic models. First, games are "essentially a one-shot-affair," but many real-life conflicts require participants to consider the ramifications of a sequence of actions. Simply selecting the best option available at each point of a conflict (the greedy approach) can have detrimental consequences for a player in subsequent stages of play [54]. Second, games assume players know all strategies and payoff functions of the game. In real-life conflicts, participants usually have only partial knowledge of the strategies available, and it may be impossible to determine actual payoffs.

These criticisms motivated several extensions to game theory—extensions which have enabled researchers to model deception. They can be broadly classified as being deception games, Bayesian games, signaling games (a type of Bayesian game), or belonging to hypergame theory. These three game classes are described below. Several other classes are used in the literature, but these do not fall into one of the aforementioned categories. In Section 2.5, they are presented in the *Other Models* subsection.

One feature shared by all classes is the rules of the game limit the players' perception of the conflict: they are games with incomplete information, imperfect information, or both. Since incomplete and imperfect information are often confused, the following subsections provide a brief review.

Imperfect Information.

In games with *imperfect information* at least one player is unaware of the actions chosen by the other players, and this player is unable to distinguish between which of several states the game could be in after the other players have taken these hidden actions. In extensive form games, nodes (i.e. the game states) are partitioned into information sets, and a player with imperfect information cannot distinguish between nodes in the same information set. In contrast, a game with perfect information is one where all states are known, i.e. all information sets are singletons. Figure 1 shows an extensive form game where state A and state B are contained in a single

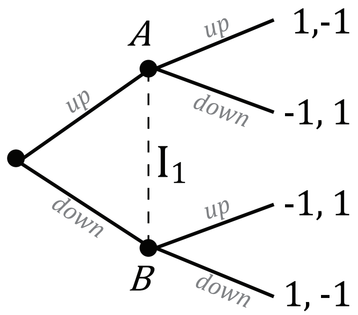

Figure 1. Example extensive form game with imperfect information. Play begins at the leftmost decion node. The first player selects a direction–either up or down. The second player cannot distinguish between state A and state B because they are both members the information set I_1.

information set (indicated by the dashed line in the figure). The first player selects an action—either up or down. The second player subsequently chooses to go either up or down, but the second player cannot distinguish whether or not the game is in state A or state B when making the decision—the two states are elements of the same information set and, and he was not provided information about the first player's choice.

Kriegspiel Chess is a board game with imperfect information. It is similar to a combination of battleship and chess. The players setup their own pieces on separate game boards, hidden from one another. A referee observes the actions of both players and tracks the moves on a third (master) board which is also hidden from the players. Play proceeds as in a standard game of chess, except that the players do not know what moves have been made by their opponent and are unable to determine the board state immediately after the opponent has moved. On their turn, the player attempts a move. If the move is valid, the referee records it on the master board without announcing what move was made. Then play is passed to the next player. If the move is invalid, the referee requests a new move from the moving player.

Incomplete Information.

A game has *incomplete information* if one or more players lacks some information about their opponents, e.g. their type, available strategies, payoffs, preferences or some combination thereof. Incomplete information games are sometimes called partial information games [37, 38]. Perhaps the most common type of incomplete information is when a player is unaware of the opponents' payouts. Incomplete information arises in many real-world conflicts, and games that include incomplete information are numerous. Negotiations can be modeled as a game of incomplete information: each participant does not usually know how much the other party values each of the different possible outcomes. In cyber warfare, attackers use the fact that software vendors do not know all possible attack vectors when new software is released. An attacker can take advantage of these vulnerabilities before they can be addressed . So, the conflict between an attacker and vendor can be modeled as a game of incomplete information, with the vendor not knowing the full set of actions available to the attacker.

Other models.

Although most games are described as having either incomplete or imperfect information, two less-common terms are encountered in the literature: informational asymmetry and interrupted observation. They are mentioned briefly for completeness. Informational asymmetry [1] arises in situations where one player has more or better information than the others. Interrupted observation indicates that a player is only able to observe the true state of the conflict intermittently. For example, interrupted observation is used in a pursuit-evasion differential game [90] where one player—the pursuer—intermittently observes the opponent; the pursuer must extrapolate the evader's new location based on its last known location and velocity.

Models of Deception.

Three classes of games constitute the majority of the deception literature. They are Deception Games [4], Bayesian Games [36], and a special kind of Bayesian Game known as a Signaling Game [20]. Several studies use Hypergame Theory [7]—an extension of game theory. The three classes named above and Hypergame Theory are described.

Deception Games.

A Deception Game is a two-player, zero-sum game where information in a vector is distorted by the deceiver. Deception Games are less frequently studied, but they have been effectively used as models for deception in political negotiations [91] and network security [61]. The first formal Deception Game was introduced by Mark Thompson in his Harvard undergraduate thesis in 1970 and later posed as an open problem in [73]. A general deception game was defined in 1988 by Baston and Bostock [4]. The general form involves two players and an n-tuple $X := (x_1, \ldots, x_n)$ of independent and identically distributed (IID) real-valued random variables obtained according to probability function P_i on the closed interval $[0, 1] \subset \mathbb{R}$. At the start of the game, Player 2—the deceiver—observes X, and presents a modified n-tuple, $Y := (y_1, \ldots, y_n)$, to Player 1. The misrepresentation, Y, satisfies the condition that at most $n - k \geq 0$ positions of X have been changed, i.e. $x_i = y_i$ for at least $k \leq n$ indices i. Player 2 presents Y to Player 1, who subsequently selects a single position $j \in \{1, \ldots, n\}$. The payout is determined by the value of position j in the original game, i.e. $(x_j, -x_j)$. The game is solved in [51], showing that for $n = 3$ and $k = 1$ the information can be completely distorted, but with $n = 4$, the optimal strategy for mark has an expected gain greater than the mean value of a random selection strategy. The generalized form of the game is solved in [30], showing for $k > n/2$, the optimal deceptive strategy completely distorts the information, rendering it useless to the mark.

The general Deception Game is single-sided, i.e. there is only one deceiving player. The related Cover-up Game [5] is extended in [68] to allow two-sided deception. The two-sided Cover-up Game is given as follows: Let $X := (x_1, x_2)$ and $Y := (y_1, y_2)$ be IID random variables selected from a probability distribution over $[0, 1] \subset \mathbb{R}$. Player 1 observes X and selects a critical value $\theta_1 \in [0, 1]$. Likewise, Player 2 observes Y and selects a critical value $\theta_2 \in [0, 1]$. Player 1 (2) reveals $x_i \in X$ ($y_i \in Y$) that is closest to θ_1 (θ_2). If a player's critical value is greater than the opponent's revealed number, the player receives as his payout the value of the covered-up number. Otherwise, the player's critical value is less than the opponent's revealed number. So, the player receives a payout equal to the value of the revealed number. To illustrate: Let $X = (0.25, 0.75)$ and $Y = (0.33, 0.66)$, and suppose the players select critical values $\theta_1 = 0.45$ and $\theta_2 = 0.94$. Then Player 1 must reveal $x_1 = 0.25$ because it is the closest $x_i \in X$ to θ_1, and Player 2 must reveal $y_2 = 0.66$ because it is the closest $y_i \in Y$ to θ_2. Since $y_2 = 0.66 > 0.45 = \theta_1$, Player 1 receives the value of Player 2's covered number, i.e. 0.33. And since $x_1 = 0.25 < 0.94 = \theta_2$, Player 2 receives Player 1's covered-up value, 0.75.

Bayesian Games.

Bayesian games were introduced in 1968 to address the issue of player misperception [36]. A Bayesian game accounts for the incompleteness of information by using probability distributions (called priors) to represent the player perceptions. Priors indicate the perceived likelihood that each player is of a particular type. In this way, a Bayesian game transforms an incomplete information game into a complete information game by augmenting the model with a set of types and a prior for the players' subjective view of the conflict. Thus, a Bayesian game is a complete information game: the players are fully aware of the type sets and each type's subjective prior.

Signaling Games.

A signaling game [20] is a particular kind Bayesian Game involving a sender (S) and receiver (R). A third player—Nature (N)—is introduced that assigns S a type according to a common-knowledge probability distribution. After observing its type t, S takes an action—literally, S transmits a signal. R selects an action after observing the signal from S. Formally, a signaling game can be specified as a tuple $\Gamma = \langle T, M, A, U_S, U_R \rangle$, where $T = \{t_1, \ldots, t_I\}$ is the set of types held by the sender, $M = \{m_1, \ldots, m_I\}$ is the set of sender messages, $A = a_1, \ldots, a_k$ is the set of receiver actions, and $U_s(t_i, m_j, a_k)$ and $U_R(t_i, m_j, a_k)$ are the utility functions of the sender and receiver, respectively.

The first signaling game [20] describes a two-player game called Wimp-or-Surly. The situation is depicted graphically in Figure 2. In this game, Player A is either a wimp or is surly, and A must choose what to eat for breakfast (either beer or quiche). Player B must decide to either duel or not duel A. Since, B does not know A's actual disposition; B must infer it based only on the signal sent by A, e.g. what A ate for breakfast. Although quiche benefits a wimp, it is more likely to signal to B that A is a wimp and encourage the duel (regardless of him being a wimp or not). On the other hand, the game says beer only benefits a surly, and since anyone can drink beer, the signal is not unambiguous. Thus, A chooses a meal after observing his own type, and B chooses to either duel or not duel. Payouts are awarded according to each player's decision.

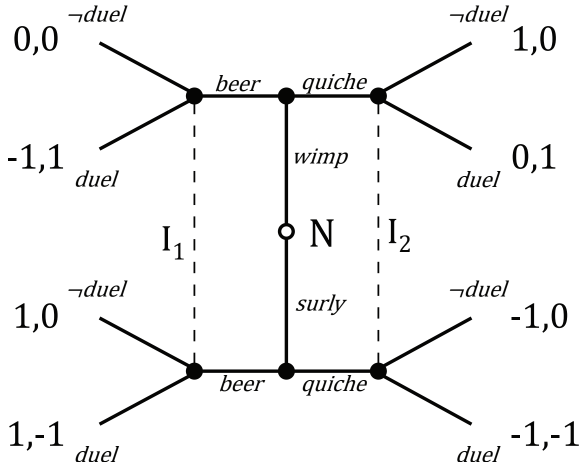

Figure 2. Wimp or Surly: the first Signaling game. Play begins when Nature (N) selects a type t for A, either wimp or surly. At the next decision node, A selects either beer or quiche as the signal (S). B observes the signal, but the information sets I_1 and I_2 make it impossible to discern t. So, B selects an action knowing only what A ate for breakfast.

Hypergame Theory.

Game theory traditionally operates under the assumption that players are well-informed about the game being played: The players are aware of the range of strategies and their opponents' preferences. However, in many real-world circumstances, this is not the case; players' perceptions of the situation and of their opponents are inconsistent. In fact, players may not know that particular actions are possible or even that certain other players exist. Even in simple conflicts, actions by one participant can appear irrational when observed by a participant with a different view of the conflict. But games of incomplete and imperfect information do not address situations where players are absolutely convinced of something that is inconsistent with reality, i.e. situations of paranoia or self-deception [32]. A deceiving player may take advantage of these false perceptions. Indeed, the best deception is often the one which causes the mark to be absolutely certain of the wrong thing.

Hypergame Theory is introduced in [7] to address situations where player perceptions did not coincide. A simple n-player *Hypergame* is defined as follows:

H.1 A set P_n of n elements, interpreted as the *players* of the hypergame,

H.2 For each $p, q \in P_n$, a non-empty finite set $S_p^q \subset \mathbb{N}$ interpreted as *p's available strategy set according to q*

H.3 For each $p, q \in P_n$, an ordering relationship O_p^q defined over the product space $S_1^q \times \cdots \times S_n^q$, interpreted as *p's outcome preferences according to q*

Hereafter, this formulation is called the *Bennett hypergame.* In this formulation, each subgame G_p^q $(= \langle S_p^q, O_p^q \rangle)$ represents q's perception of the game that player p is playing. So, the collection of sets $\{S_1^q, \ldots, S_n^q\}$ together with $\{O_1^q, \ldots, O_n^q\}$ represent q's perception of the hypergame, i.e. the collection G^q $(=\{G_1^q, \ldots, G_n^q\})$ describes the hypergame that player q is playing. So, the hypergame $H = \{G^1, \ldots, G^n\}$ is interpreted as a set of n games, each expressing a different player's perspective.

Bennett hypergames deviate from the original utility value treatment of von Neumann and Morgenstern [83] by using preference vectors to represent player preferences for each outcome. A preference vector is simply an ordering over outcomes. When using a preference vector in games, the outcomes are first arbitrarily assigned an index. Each player then sorts the outcome indices according to their preference: the most desirable outcomes are listed first and the least desirable are listed last. The ordinal approach reflects the fact that Hypergame Theory was originally introduced as a "soft" approach to conflict analysis [10]—i.e. to help structure and clarify the situation for decision-makers.

An example Bennett hypergame, H, is loosely based on the Battle of the Sexes. Here $X \succ Y$ indicates a preference for X over Y, and $X \sim Y$ indicates indifference between X and Y.

- $P_n = \{1, 2\}$

- $S_1^1 = \{a, b\}, S_2^1 = \{A, B\}$

  $S_1^2 = \{a, b\}$, $S_2^2 = \{A, B, C\}$

- $O_1^1 = O_1^2 = \{(a, A) \succ (b, B) \succ (a, B) \sim (b, A)\}$

  $O_2^1 = \{(b, B) \succ (a, A) \succ (a, B) \sim (b, A)\}$

  $O_2^2 = \{(a, C) \sim (b, C) \succ (b, B) \succ (a, A) \succ (a, B) \sim (b, A)\}$

Although originally introduced using ordinal preference vectors, hypergames allow for cardinal utilities which results in a special case of the ordinal system above [6]. The hypergame in Figure 3 is equivalent to H except that the outcome preferences have been assigned cardinal values. H is represented as three games in normal form. In this example, Player 1 believes that both players are playing an instance of *Battle of the Sexes*: the players are trying to decide whether to see movie A or movie B. If the players go to different movies, then they gain nothing. Player 1 prefers movie A to movie B, and Player 2 prefers movie B to movie A. Player 2 accurately perceives Player 1's game (since $G_1^1 = G_2^1 = G_1^2$); however, Player 2 is also aware of a third strategy, see movie C, which is preferred regardless of whether Player 1 attends or not. Player 2's rational course of action is to play the pure strategy C, while Player 1 mixes over A and B according to the mixed strategy Nash equilibrium $(A, B) = (\frac{2}{3}, \frac{1}{3})$. This example demonstrates how players in a hypergame need not agree upon the game being played.

Fraser and Hipel introduce a slightly modified formulation that exclusively considers preference vectors [26]. The preference vector for player i is an ordering over the outcomes that expresses each outcome's desirability according to player i. The preference vectors are sorted from most-desirable outcome to least-desirable outcome. A Fraser-Hipel hypergame is defined entirely by the preference vectors of the players and their beliefs about each other's preferences vector. Thus, player q's game is

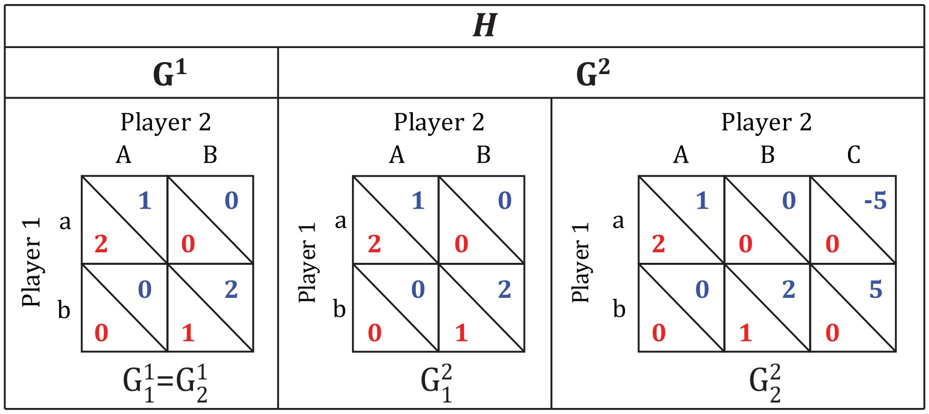

Figure 3. Example Hypergame. Player 1 perceives the game as an instance of Battle of the Sexes. Player 2 accurately perceives Player 1's game, but is aware of a third strategy, C, that is hidden from Player 1.

defined as $G^q := (V_1^q, V_2^q, \ldots, V_n^q)$, where V_i^q is player i's preference vector according to player q. An n-player Fraser-Hipel hypergame is therefore defined by the tuple of subgames as $H := (G^1, G^2, \ldots, G^n)$. Takahashi, Fraser and Hipel present several algorithmic approaches to analyzing stability in Fraser-Hipel hypergames [26,75], and the form is used to model several complex conflicts [27–29].

Since its introduction, hypergame theory has been used to study real-world conflicts in business, sports, resource allocation problems, and in the cyber domain. Hypergames were also used to analyze several military conflicts including the 1940 fall of France [9], the 1956 Suez Canal Crisis [70], the 1973 Middle East war [67], and the 1982 Falkand–Malvinas conflict between Argentina and Britain [41]. These *post-hoc* analyses demonstrated the viability of Hypergames as an effective tool for inference, i.e. they can help researchers understand the reason for a particular conflict's outcome.

The first live test of hypergame theory was performed by Bennett *et al.* during the worldwide shipping crisis of the 1970s and 1980s [11]. The crisis is attributed (primarily) to over-production and under-demand of large shipping vessels and a general unwillingness of the participants to adapt to the changing economic environment. in 1978, Bennett *et al.* published the a report on their hypergame analysis. They predicted an unexciting outcome: no involved party had an incentive to unilaterally deviate from their current strategy. Thus, they anticipated the shipping crisis to continue without any significant changes. The authors revisited the topic two years after publishing their initial report to summarize the state of the crisis. Their retrospect evaluation showed that hypergame theory did provide a means of predicting conflict outcomes; the crisis continued as predicted to the detriment of many parties involved. In 1991, the Persian Gulf crisis between Iraq and the United States-led Allied forces provided another opportunity to apply the predictive power of hypergames to a military conflict. On January 11 and 12, 1991, Wang and Hipel [87] modeled and analyzed the Persian Gulf war as a 1st- and 2nd-level hypergame just before the outbreak of the air campaign on January 16, 1991. They predicted U.S.-led forces would launch air strikes against strategic targets and later conduct a full-scale ground war, while Iraq would respond with military offensives including Scud missiles, ground battles, and the use of non-conventional weapons. Their predictions coincide well with the historical outcome and demonstrated the viability of hypergame theory as a tool for predicting the outcome of complex conflicts with misperception or deception.

2.3 The Practicality of Deception

Game theory has been used to determine if a situation warrants the use of deception, i.e. benefit the deceiver. From a game theoretic perspective, the desirability of an outcome is encapsulated in the outcome's payout value. So, for a deception to benefit the deceiver, it must necessarily improve the expected payout. But this is not the only measure of a deception's desirability. Game theorists have further characterized a deceptive strategy according to other features, e.g. surety [33] and stealth [16,33]. Thus, game theory provides a rich set of tools for modeling conflicts with deception and misperception. This section discusses how game theory is used to determine whether a situation warrants the use of deception. It also introduces several desirable properties and the five GT mechanisms for performing deception.

Situations that Warrant the use of Deception.

One aspect of deception planning is to distinguish situations that warrant the use of deception from those that do not. Drawing inspiration from interdependence theory (a subfield of social exchange theory formalized by Thibaut and Kelley [76]), Wagner and Arkin [84] use outcome correspondence and payout interdependence to measure the degree to which a game's payouts warrant the use of deception.

Interdependence describes the extent to which the actions of one player impact the reward of another [86]. Interdependence is calculated separately for each player and ranges from 0 for independent situations to +1 for dependent situations [84]. When payouts are independent, actions by one player have no impact on other players' utility. The players can act as if no other player exists. Although game theory can be used to analyze conflicts without payout interdependence, other approaches are more appropriate, e.g. decision theory [60].

"Correspondence describes the extent to which the outcomes of one individual in a situation are consistent with the outcomes of the other individual" [86]. When correspondence is high, players select mutually beneficial actions to maximize their own utility. When correspondence is low, one player gains as the other player(s) lose. So, games with low correspondence favor competition over cooperation.

Wagner and Arkin [85] argue that deception is only warranted when correspondence is low: if correspondence is high, a non-deceptive (honest) signal can be used to improve each player's outcome. This does not imply that deception *cannot* be

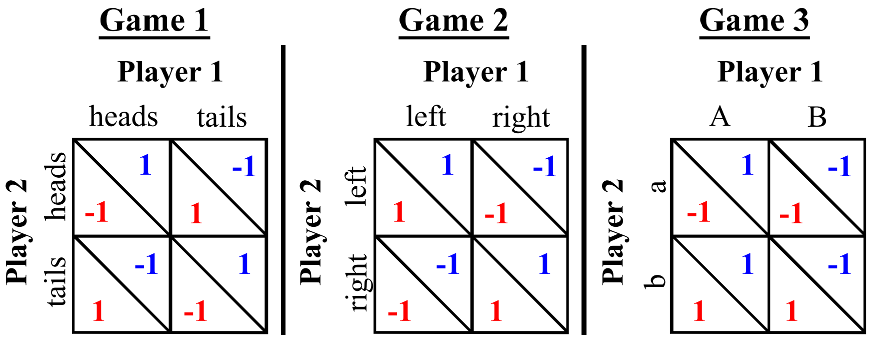

Game 1

Player 2 \ Player 1	heads	tails
heads	1 -1	-1 1
tails	-1 1	1 -1

Game 2

Player 2 \ Player 1	left	right
left	1 1	-1 -1
right	-1 -1	1 1

Game 3

Player 2 \ Player 1	A	B
a	1 -1	-1 -1
b	1 1	-1 1

Figure 4. Payout interdependence and correspondence. Three games illustrate payout interdependence and correspondence. Game 1 (left) is a non-cooperative game with interdependence and low payouts correspondence. Game 2 (center) is a cooperative game with interdependence and high payout correspondence. Game 3 (right) is not interdependent, not corresponding, and not necessarily cooperative or non-cooperative; it could be analyzed more simply using decision theory.

used when correspondence is high, but it does imply that an honest signal could be used instead. For example, in Battle of the Sexes, two players must coordinate their actions: if they select the same action, they win; if they select different actions, they lose. The players could rely on a public signal such as a traffic light or the weather to coordinate their actions. In this case, they can improve their expected outcome using a non-deceptive (honest) signal. However, the signal need not be honest. A deceptive, false signal can be used to coordinate their actions, as well. So, even though Battle of the Sexes has high correspondence, both an honest signal and a deceptive signal can improve the players' outcomes.

It is important to make a distinction: A game may exhibit interdependence between payouts, but that does not necessarily mean it is either cooperative or non-cooperative. To help illustrate this difference, consider the three games in Figure 4. Game 1 is an instance of matching pennies: it is an example of a non-cooperative game with interdependence and conflict in the payouts. Game 2 is an instance of the Choosing Game where the players win if they choose the same side and lose oth-

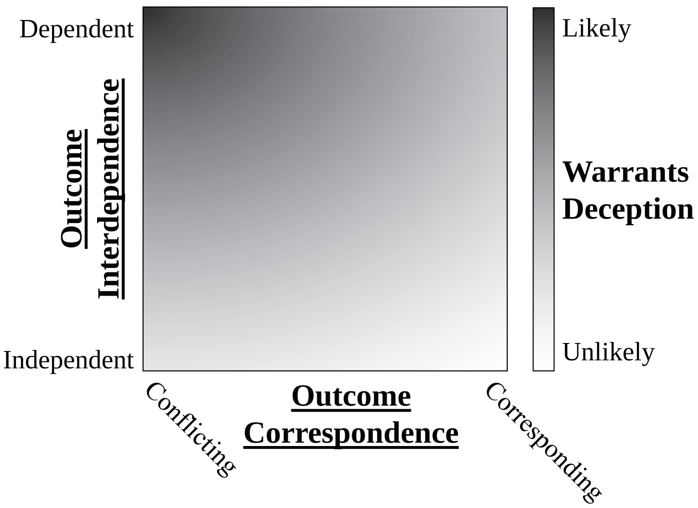

Figure 5. Two-Dimensional Correspondence Space. A two-dimensional representation of interdependence space [85] showing outcome correspondence on the horizontal axis and outcome interdependence on the vertical axis. Games that map to the dark region tend to warrant the use of deception more than games mapped to the light region.

erwise. The Choosing Game is cooperative with interdependence, but it does not exhibit conflict in the payouts. Finally, Game 3 is a contrived example that does not exhibit interdependence, and its payouts do not conflict; it is neither cooperative nor non-cooperative, and deception is entirely unwarranted.

A game is mapped into a two-dimensional space according to its outcome interdependence and outcome correspondence to determine the level to which it warrants the use of deception. This space is illustrated in Figure 5. Games that map to dark regions tend warrant the use of deception more than games mapped into the light regions. The upper-left corner consists of games that maximally warrant the use of deception. Areas on the lower area tend not to warrant the use of deception because the actions of the mark have a lesser impact on the outcome of the deceiver. Games with high correspondence do not require false communication to improve the deceiver's outcome, and therefore deception tends not to be warranted.

Necessary Conditions for Deception.

To perform deception, two conditions must be satisfied. First, the deceiver must have the capacity to deceive. Second, the conflict must be one of incomplete or imperfect information. This section describes the necessary conditions for deception.

Capacity to Deceive.

The first condition is that deceptive player be able to influence and take advantage of the mark's (mis)perceptions, i.e. that the deceivers have the capacity to deceive. This condition may be unsatisfied for several reasons including resource and time constraints, cognitive ability, or when player actions are restricted through the course of play.

For example, several articles consider a hypothetical conflict between network attacker and network defender [18,31,61]. The defender places camouflaged monitoring systems—called honeypots—throughout the network to lure attackers away from critical systems. In [18], the attacker can probe a system to distinguish honeypot from non-honeypot before deciding whether or not to attack. For an attacker with unlimited resources, one solution is to simply probe every system before attacking. Since the probe eventually returns the true system type, the attacker can always distinguish normal systems from honeypots. In this case, the network defender has no capacity to deceive unless time is an important consideration for the attacker.

Incomplete or Imperfect Information.

The second condition necessary for deception requires the mark to misperceive some aspect of the conflict. Without complete or perfect knowledge, each participant acts based on a limited perception of the conflict. Over time, the participants may update their perception based on the past experiences and observations. However, sometimes, a deception is designed to take advantage of learning on the part of the mark. For instance, Burns [17] describes Operation Overlord where the allied forces of World War II performed an extensive campaign of deception surrounding the place and time of their Normandy invasion. Their efforts so deceived the defending Germans of their intentions, that the German 15th Army delayed their response for several weeks expecting the true offensive to begin several hundred miles to the North near the Pas de Calais.

In games with uncertainty, learning from past mistakes can help players identify inconsistencies in their perception of the environment. In Bayesian games, players learn by updating their beliefs about an opponent's type according to Bayes' rule, and these beliefs may change on the basis of their actions. However, the ability to learn does not necessarily imply the mark's perceptions converge to toward reality [32]. It only implies that the mark's perceptions are updated. Some deceptive techniques take advantage of a player's ability to learn. To do this, the deceiver carefully plans actions to precondition the mark. As the mark learns from past interactions with the deceiver, his perception of the conflict tends to diverge from the true nature of the conflict. Thus, if the information revealed to the mark is carefully planned, the deceiver can mislead the mark in an advantageous way. For example, Gharesifard and Cortés [32] introduce a method for learning opponent preferences in hypergames where players interact multiple times. They later introduce an algorithm [33] that can be used to exploit learning on the part of the mark through carefully planned action sequences.

Desirable Properties.

Given that the necessary conditions for the deception are met, it is useful to also evaluate the desirability of the a deceptive strategy. An natural measure of desirability is the expected value gain, i.e. the expected value of the deceptive strategy minus the expected value of the game at it's true equilibrium. Although this technique is effective, researchers have expanded the way in which the quality of a deception can be measured by identifying several other desirable properties. A description of three such properties, namely Stealth, Surety, and Cost are described. However, an important property that is absent from the literature is risk of exposure, i.e. the possible value loss for being out-witted by a deception-aware mark. If the deceptive strategy requires off-equilibrium play on the part of the deceiver, then the deceiver has the potential to lose value. Addressing risk in game theoretic models is important because of the central role that risk analysis plays in real-life deceptive planning. Risk must be reexamined in every stage of the deception [77] because failure or disclosure of the deception can significantly impact the outcome.

Stealthy.

A deception is *stealthy* if the deceiver purposefully restricts his actions to those which do not contradict the mark's beliefs. Stealthy deceptions are possible in games with asymmetric information—when the deceiver has more or better information than the mark. The deceiver is able to leverage this information superiority to avoid strategies that reveal helpful information to the mark. A stealthy deception is played until the deceiver is confident that the desired outcome is inevitable. At that moment, the deceiver may choose to either continue restricting his actions or not. If the deceiver continues to play restricted actions, the stealthy deception is called *tacit*. Otherwise, the deception is said to be *revealed*. This terminology is introduced in the Voting Game [16] (described below). When the deceiver misrepresents his outcome preferences, the other players respond in a rational way. The deceiver can take advantage of the other players' responses in two ways: He can tacitly deceive and receive his second-preferred outcome, or he can play the revealed strategy and receive his most-preferred outcome. The choice depends on the number of interactions. A tacit deception is especially desirable in repeated conflicts because it can succeed repeatedly without being revealed to mark. The long-term benefits of a tacit deception may outweigh the gains of a deception that is revealed immediately.

Surety (Certainty).

Surety is a property of the outcomes in a game with sequential actions. Surety indicates whether or not a action sequence exists that guarantees an outcome from the game's current state. Formally, let S be the set of game states and let O be the set of outcomes. An outcome $o \in O$ is surely deceivable if there exists a deceptive strategy for the deceiver from the game's current state $s \in S$ that achieves o with probability one. Suppose the deceiver desires outcome o and that o is surely deceivable from state s. Then, in a way, the game ends once it reaches state s; from that moment forward, the players simply go through the motions to reach outcome o. When $o \in O$ is surely deceivable regardless of the initial state (i.e. $\forall s_k \in S$), then o is said to be surely strong deceivable. The definitions for surely deceivable and surely strong deceivable are given in the context of hypergames in [33], but these concepts can be used to describe deception in general.

Cost.

The cost of deception describes the amount of effort that must be expended to achieve the desired effect. Cost can represent many different aspects of a conflict, e.g. profit loss, loss of life, opportunity cost, reputation damage, regret. Surprisingly, most studies never explicitly address the cost of deception. Instead, they assume a zero-cost deception. Future work might take a different approach by parameterizing the payouts based on the cost of performing the deception. This approach is often used in Costly Signaling Theory to study cooperative communication. For example, Floreano *et al.* [25] studied the evolutionary conditions for the emergence of communication. After simulating the evolution of 100 colonies of 10 robots over 500 generations, the authors implement the resulting communication strategies in physical robots. The study presents three major findings. First, deceptive strategies emerge more often among unrelated robots, i.e. those randomly grouped together from dissimilar evolutionary colonies. Second, deception reduces the overall productivity of the group compared to those groups that did not use deception. Third, deception is rational even when deceptive signaling is costly.

Traditionally, game theory incorporates all aspects of value to the player in their payouts: e.g. it represents an outcome's value as a scalar equal to the reward minus the cost. In many conflicts, this approach is appropriate: if the cost of performing an action is $100 and the associated reward is $1000, the outcome's payout is set to $900. But in many real-world conflicts, the reward and the cost are not comparable. How does one subtract regret from profit margin? It is conceivable that one could assign a monetary value to the regret (e.g. hours of therapy multiplied by therapist cost per hour), but a more direct approach considers cost as a separate dimension from the payout itself. The game theorist can then take a multi-objective approach to the problem of performing deception: maximize the expected value gain while simultaneously minimizing the cost of performing deception.

Deception Mechanism.

A deceptive mechanism describes the means by which the deception is accomplished. From a game theoretic perspective, the mechanism targets the mark's perception of the game, e.g. payouts, action set, players set, player types, information set. Surprisingly, the surveyed articles each consider at most a single deceptive mechanism. The remainder of this section describes the various deceptive mechanisms in game theoretic terms.

Environmental Deception (Payout Manipulation).

Environmental Deception (or payout manipulation) is a mechanism that causes the mark to misperceive the payouts of the game. Since the desirability of an outcome depends on the payout, environmental deception causes the mark to select suboptimal strategies and allows the deceiver to improve his outcome. Camouflage can be used to perform environmental deception. Camouflage causes the mark to misperceive the threat posed by (or the value of) the camouflaged object.

This mechanism naturally leads to the question: How much must a payout value be manipulated before the game's equilibrium is changed? A strategy is called *essential* if and only if it has a non-zero probability at equilibrium. Arsham [2] presents a necessary and sufficient condition for such strategies to be stable, i.e. to remain essential when payoffs are changed. This is relevant to environmental deception because the deceiver is often trying to induce changes to the opponents' strategy. The condition can be used to determine the degree to which an environmental deception may alter the payouts before an action enters/exits the essential strategy set. If environmental deception is costly (vice cost-free), this approach can help identify least-cost deceptions.

Tactical Deception.

A tactical deception occurs when an honest action is used out-of-context so that others misinterpret its meaning [88]. Tactical deception is often achieved in real-world conflicts by withholding information. For example, if a person who usually plays golf on Saturdays secretly uses the time instead to prepare for their upcoming anniversary, they perform a tactical deception. This mechanism is not exclusive to humans; it is observed in many animal species, as well [88].

In the context of a game, a tactical deception uses or manipulates the mark's perception of the information sets to the deceiver's advantage. In effect, a tactical deception prevents the mark from distinguishing between states. Many games use information sets to model uncertainty and deception, e.g. in Signaling Games, but no example could be found of a deception that allows changes to the information sets.

To illustrate such a situation, consider the conflict depicted in Figure 6. Here, the attacker is the deceiver, and the defender is the mark. Suppose the defender must guard two targets, t_1 and t_2. The defender controls a single unit, which he can send to either of the two targets. The defender also has an early-warning system that announces the destination of any incoming attackers. If the defender intercepts the attacker, it receives a payout of one. Otherwise, the defender receives a payout of negative one. The payouts in this game are zero-sum. Suppose an attacker seeks to destroy one of the targets and has developed a jamming capability that can disrupt the early-warning system. In effect, this prevents the defender from being able to distinguish the state where the attacker strikes t_1 from the state where the attacker strikes t_2.

Action Deception.

An action deception occurs when the deceiver's causes the mark to misperceive the action set. In normal form games, an action deception hides actions from the mark. Vane [79] shows that when the mark plays a subset of the available actions, the best response is often a pure strategy. In extensive form games on directed graphs, an action deception eliminates edges from the digraph. If action deception can be used

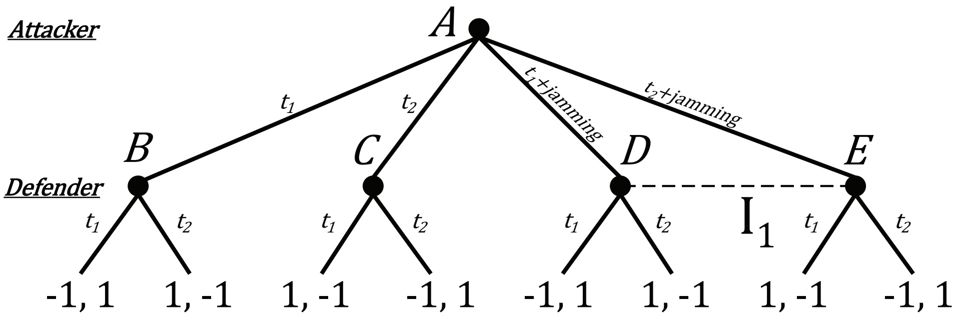

Figure 6. Example of Tactical Deception. The attacker selects target t_1 or t_2 to attack and can employ jamming. If jammed, the defender cannot distinguish between states D and E because they are both members of information set I_1.

on an edge, then the edge is said to be edge deceivable. If properly employed, an action deception causes the mark to select a sequence of actions that would otherwise be irrational.

Gharesifard and Cortes [33] study action deceptions in hypergames using the H-digraph [32]. The H-digraph is a directed graph. The vertices correspond to the payouts of each outcome and the edges correspond the actions available to each player. The authors use the order of each player's preference lists to determine when an outcome is an improvement—if outcome y is preferred to outcome x by player i (denoted $x \preceq_i y$), and all other players are indifferent, then y is said to be an improvement to x. The turn-based nature of their game causes some outcomes to be sanctioned by a player. An outcome x is sanctioned by an outcome y according to player i if either (a) y is preferred to x by i and all other players are strategically indifferent to the two, or (b) if player i is strategically indifferent between x and y but y is an improvement for all other players. An outcome is rational if it cannot be improved upon, and an outcome is *sequentially rational* if it cannot be improved to a sanction-free outcome. Thus, no rational sequence of actions can attain a sanctioned outcome. Using the notation presented in the first several sections, the authors present a necessary and a sufficient condition for determining whether action deception can be used to remove an edge from the H-digraph. Furthermore, they “fully characterize when [stealthy] deception is possible” and present a algorithm to “find a sequence of deceiving actions” [33].

Participant Deception (Hidden/Fictional Players).

Another deceptive mechanism is to influence the mark’s perception of the set of players, i.e. to either hide the existence of a real player or to introduce fictional players. The latter version could occur in conflict at school: one child invokes the threat of his (fictional) big brother to avoid being teased by other children. The former version could occur in a conflict between nations, where one is supported secretly by an outside entity and provokes their opponents. Both hypergame theory and Bayesian games can be used to model conflicts where the existence of one or more players is unknown to some or all of the other players. Sasaki [69] argues that despite the fact that the two modeling techniques are equivalent, hypergames may be better suited for this type of analysis: hypergames tend to have a more natural and less ambiguous interpretation than Bayesian games: Bayesian games assume that the set of players is common knowledge, but in hypergames, players have their own subjective view of the entire conflict, including the set of players.

Misrepresenting Player Type.

The final mechanism for deception is to distort the mark’s perception of the players’ types. The misrepresentation mechanism requires that the mark respond differently based on the player types and is the key mechanism for deception in Signaling Games: The player is assigned a type by nature, and the signal they send communicates something about their type.

This mechanism is common in many real-world conflicts. One example of this mechanism is a hustle—a scheme in which the mark is deceived after the opponent first establishes a history of play as an especially weak type when the potential for loss is low. Later when the potential for gain is high, the deceiver capitalizes on the mark's misperception by playing according to his true type and surprising the mark. This is a common approach in poker where part of the challenge is to discern the other player's type—whether they play only strong hands or bluff on weak hands too, whether they check and call frequently or else bet and raise frequently. A poker player may misrepresent his type early through slow play [71, 72] to gain an advantage over his opponents when the stakes are higher.

2.4 Detecting Deception

Deception detection is studied in many fields, but little has been done from a game theoretic perspective. This is surprising considering the central role that deception detection plays in counterdeception. By definition, counterdeception requires the ability to detect deception i.e. it assumes the player knows his opponent is attempting to deceive. "Decision makers must be aware of adversary deception activities so they can formulate informed and coordinated responses" [77]. However, deception detection is helpful for the deceiver, as well: knowing the indicators that expose a deceptive strategy, a deceiver can maintain the deception long by avoiding compromising or revealing actions. For example, after the Allied code breakers in World War II cracked the Enigma code, they were able to decode the Nazi transmissions. The Allies concealed this achievement so that the decrypted information could be exploited in a strategic way. If the Allies were careless in using the information they obtained, the Nazis would have realized their communications were compromised and adopted an alternative encryption scheme. Thus, the Allies had to select only a subset of information upon which they would take action. This example illustrates how knowing which indicators expose a deceptive strategy can enhance the long-term maintainability of the deception.

Several studies [14, 16, 86] consider the impact of deception against deception-naïve opponents. Future work could extend these efforts, studying deception against deception-aware opponents. This is an area that requires additional investigation, especially in the context of repeated games, where multiple interactions provide an opportunity to reveal a deception. Hoang*et al.* [42] survey repeated games in network security. They describe several studies in cooperative games where nodes in a network are either good or bad. The nodes must differentiate the good from the bad. Although the bad nodes do not necessarily employ deception, these studies give credence to the idea that game theory can be used to differentiate honest and dishonest actors in a network simply by observing their behavior. Much work has been done to detect deception in other fields. For instance, Elsaesser and Stech [23] describe a process to assist intelligence analysts in deception detection based on an analysis of competing hypothesis [40]. In hypergames, an algorithm was recently developed to identify and exploit inconsistencies in player perceptions [33]. The algorithm makes two assumptions. First, it assumes the deceiver has perfect knowledge of the mark's game. Second, it assumes that if the deceiver takes actions that contradict the mark's perceptions, the mark will update his perception of the game accordingly. Future work might explore the situation when the deceiver does not have perfect knowledge of the mark's game a priori [43]. In this case, it is possible that a deception-wary mark might benefit from intentionally playing a suboptimal strategy. These *exploratory actions* could prevent the deceiver from being able to anticipate the mark's actions perfectly, make deception more difficult, and may uncover a method for deception-aware players to detect an act of deception.

2.5 Performing Deception

This section describes the various studies focused on performing deception. This section is divided into four subsections according to the game type. The first three subsections describe deception in Deception Games, Hypergame Theory, and Bayesian Games, respectively. The last subsection describes how various other game types were used to model deception.

In Deception Games.

Deception Games are useful especially when the information being distorted represents player preferences. Misrepresenting preferences is studied in two-player, 2×2 games [16] and in simple three-player voting games [14]. This approach is used to explain the actions of the United States at the Geneva Conference of 1954—the conference which led to the partitioning of Vietnam into the northern and southern regions [91]. The voting game consists of three voters (v_1, v_2, v_3). Each voter casts a single vote for one of three alternatives, a_1, a_2, a_3. Voter v_1 is given an additional vote in the event of a three-way tie. The player preferences are defined from most preferred outcome to least preferred outcome as follows:

- v_1: (a_1, a_2, a_3)
- v_2: (a_2, a_3, a_1)
- v_3: (a_3, a_1, a_2)

The preferences are assumed to be common knowledge among the players. Knowing that v_1 could cast a tie-breaking vote, it is not in the interest of v_2 to vote for his most-preferred outcome: it would result in a three-way tie, v_1 would cast a second vote for a_1, and v_2 would receive his least-preferred outcome. Instead, v_2 improves his expected outcome by voting for his second most-preferred outcome, a_3. Thus, despite the fact that v_1 has tiebreaking power, v_2's action guarantees v_1 receives his least-preferred outcome. It is asked if there is any recourse for v_1 through deception, i.e. whether v_1 can improve his expected outcome by misrepresenting his true preferences. The conclusion is if v_1 deceptively announced a preference of (a_2, a_1, a_3), v_2 would have an incentive to vote for a_2 and v_3 would have an incentive to vote for v_1. Thus, v_1 had the option to either perform a tacit deception by voting for a_2 (despite his true preference for a_1) or reveal his deception by voting v_1. This simple model is extended to the case when two players could deceive simultaneously [15].

Deception Games have been used to study computer network defense strategies. Specifically, the Honeypot Selection Game (HSG) [61] is based on the Deception Game and earlier work in [31]. In network security, a honeypot is a networked system designed to lure would-be attackers away from critical network resources. Since honeypots are equipped with substantial logging capabilities, they provide a way of delaying an attack and sometimes provide means of identifying the origin of the attack. The HSG provides insight into how honeypots should be allocated to maximize their effectiveness. Although the technical details of a honeypot implementation are important, the HSG highlights the strategic importance of honeypot allocation and how suboptimal allocation strategies can degrade the honeypots' effectiveness. The model is similar to the Deception Game except that rather than changing the values in the original vector X, the network defender extends X by inserting values into new positions. The rationale is simple: The network exists in the real-world and is comprised of many systems. The values in the vector X represent the value of the real-world system. By adding k honeypots to the network, the network defender is extending the network from n systems to $n + k$ systems. Thus, the resulting vector $Y = (x_1, \ldots, x_n, h_1 \ldots h_k)$ is composed of the original systems, x_i's, and the new honeypots, h_i's.

In Hypergame Theory.

Hypergame theory is used to model the 1976 negotiations between Ford Motor Company (FMC) and the governments of France, Germany, and Great Britain [8]. The negotiations were opened to determine in which country FMC would construct its newest European manufacturing plant, but early in the negotiations, it was clear to FMC that Britain was the best choice. The hypergame examines the impact of the FMC misrepresenting its preferences, playing off one bidder against another after having already secretly decided the winner. This study illustrates how such a conflict could be modeled using hypergames, how the dispenser (FMC in this case) could carry out simple forms of deception, and how inter-bidder communication can be used as a hedge against deception in similar circumstances.

Wang and Hipel [87] model the effects of Allied deception against Saddam Hussein during the first Gulf War. In the buildup ahead of the air campaign, the Allies used deception to cause the Iraqi leadership to anticipate a land invasion would originate from two places: from the Saudi-Kuwaiti border and the sea. Falling for the deception, the Iraqis focused their defenses on the border and sea, and neglected to protect their flank. However, at the onset of the air campaign, the ground forces near the Saudi-Kuwaiti border were moved secretly as far as 500km Northwest toward the central Iraqi border [87]. From this position, the Allied ground forces surprised the Iraqi leadership, avoided much of their defensive preparations, and faced little resistance.

Gharesifard and Cortes [33] provide a formal definition for an n-player, k-level hypergame with imperfect information based on earlier work by Bennet [7]. They also define improvement, rationality, sequential rationality and equilibria in such a game. The authors then present a two-player, turn-based hypergame with asymmetric information, i.e. where the deceiver has perfect information about the payouts while the mark does not. In their game, each player seeks to maximize his payout. After each action, the mark refines his estimation of the game's true payouts. When an action contradicts the mark's perception of the game, it causes a cascading update to the mark's belief structure. So, the deceiver tries to avoid such actions via stealthy deception. The key result includes a general method for computing action sequences through the hypergame's H-digraph that accomplish stealthy edge deceptions.

In Bayesian Games.

Zhuang, Bier, and Alagoz [92] examine an N-period attacker-defender resource allocation game with incomplete information. The game is converted to one of imperfect information by including a hypothetical Nature player. The game begins with Nature assigning a type to the defender according to a probability distribution known a priori to the other players. The defender selects a defense posture and a signal. The attacker observes the signal, updates his belief about the true defender type, and chooses an attack effort. The payouts are given to each player according to the defender's level of defense and the attacker's selected attack effort. This approach is novel in that the defender also has private information about the valuation of its assets. It shows that the defender benefits from secrecy and deception regarding his actual valuation of his assets.

Several works address the problem of strategic honeypot allocation to support network security. The problem is modeled a hypothetical conflict between a network attacker and network defender. The defender places camouflaged monitoring systems—called honeypots—throughout the network to lure attackers away from critical systems. Garg and Grosu [31] assume the honeypots are always successful, and Pibil, *et al.* [61] extends this model by taking a probabilistic approach. However, an experienced, patient attacker with sufficient resources (i.e. the typical modern cybercriminal) is not necessarily susceptible to such deceptions since they can repeatedly probe systems. In [18], Carroll and Grosu investigate such a situation. They use a signaling game with incomplete information to model the conflict between a network attacker and network defender, where the defender is not necessarily capable of deception. Their model depicts a heterogeneous network of normal systems and honeypots similar to [31]; however, they charge the attacker a fixed cost for determining a system's type (i.e. probing) and assume that the probes always return the true system type. Thus, their attacker could not be deceived by the defender once the system had been probed. Nonetheless, since probing is costly, the cost-constrained attacker might choose to attack a system without first probing its type. Against such an adversary, the defender has the capacity to deceive.

Other Models.

Yavin [90] studies deception in a stochastic, two-player differential game. The players represent two airplanes, one chasing the other. Both players select their strategy according to their bearing and distance: the pursuer must intercept the evader, and the evader must avoid the pursuer. Three cases are analyzed. The first case allows the evader to induce errors in the pursuer's perception of the evader's position and bearing. The second case allows the evader to produce two decoys to confuse the pursuer. The third case allows the evader to interrupt the pursuer's observations. The most interesting result is that in the third case, it is possible that a line-of-sight guidance law [53] is not an optimal pursuit strategy.

Israeli [44] shows how a player could deceive another in an infinitely repeated game by communicating several fictitious payout matrices. Essentially, the deceiver's intent is to "sow doubt" in the mark's mind as to the true state of the conflict. The true payouts for the deceiver and mark are given by matrix A and matrix B, respectively. Then the deceiver introduces a finite set of payout matrices $\{A \equiv A_1, A_2, \ldots, A_{|K|}\}$

and a probability vector $p \in \Delta^K$ that indicates the likelihood of each game being the true game. The mark must select an action based on this distorted view of the game.

Three studies [19, 37, 86] demonstrate how the availability of information can impact the success or failure of a deception. These three studies are described.

In [37], Hespanha and Ateşkan examine the use of deception on a simultaneous, zero-sum, stochastic game of attack and defense (similar to a Blotto game [12]). The game involves two players: (i) the defender guarding two locations (A and B) with three units that can be distributed between them, and (ii) the attacker who chooses a location to strike. The payouts for the defender correspond to the number of units stationed at the target selected by the attacker: c_0 if the selected target is undefended, c_1 if one unit defends, c_2 if two units defend, or c_3 if three units defend. Thus, the defender's payout is higher when the attacker strikes a well-defended location, and it is lower when the attacker strikes an undefended location. The payouts for the defender are scaled so that $c_0 = 0$, $c_3 = 1$, and $c_0 < c_1 < c_2 < c_3$. The game is zero-sum. So, it is best for the attacker to select an undefended target. For example, suppose the defender stations two units at A and one unit at B, and suppose the attacker strikes B. Since one unit was stationed at B, the defender's payout is c_1, and the attacker's payout is $-c_1$. In this case, the attacker selected the best course of action, since $-c_1 > -c_2$ and the attacker is maximizing utility.

Three versions of the game are described and analyzed. They each differ in the amount of information available to the attacker. In the first and most simple version, the attacker has no information about the distribution of the defending units. Here, the authors present a mixed strategy Nash equilibrium (MSNE). Interestingly, two cases emerge. If $c_1 + c_2 \leq 1$, then the defender plays an all-or-nothing strategy on a single location. Otherwise, the defender mixes evenly between a 2-1 and 1-2 strategy of defense. The former case occurs when the attacker is proportionately strong versus two or fewer units while the latter case implies that a single unit or pair of units can substantially attrit the attacker's forces.

The next form of the game affords the defender complete control over which units are seen. In this case, the defender can either reveal one or two units to the attacker. (Although it could reveal three, it would be to its own detriment.) The authors idealize this situation by saying that the attacker always sees a unit that is revealed. This section is useful because it provides a taxonomy of strategic policies. For instance, the attacker can either play blind (ignore detections), naïve (attack location that did not detect a unit), and counterdeception (attack location where detection occurred) policies. Likewise, the defender can play a policy of no-information (reveal nothing), deception (reveal units at the under-defended location), or disclosure (reveal unit at best-defended location).

There are several MSNEs for this version of the game, which again depend on the value of c_1 and c_2. The MSNEs imply a counter-intuitive result: even though detections provide definite information about the location of the defense units, the attacker gains nothing from them. So, the best strategy for the attacker is to either mix between naïve and counterdeception or to randomly select a location. The final and most complex game form uses a characteristic function $\chi(n_A, s_A)$ to model the conditional probability of detecting a unit at location A given that there are n_A units defending A and s_A such units shown (not camouflaged). The authors provide several practical example functions but assume a generic characteristic. The overall intent of this generalization is to account for the fact that a unit displayed may not be detected, and it is shown that the first two forms are special cases. Abstractly, the characteristic function represents a sensor's probability of detection. Reliable sensors detect units regardless of whether or not the unit is visible. Three cases arise in the analysis of this game based on the sum of c_1 and c_2 and the reliability of the sensors.

Castanon *et al.* [19] use a sequential two-player game to model a conflict between a combat aircraft and a mobile missile launch vehicle. The aircraft tries to prevent the ground vehicle from moving into position and launching a missile. However, the ground vehicle can send decoy vehicles to entice the patrolling aircraft to leave its patrol station. The vehicle strategies are to either use a decoy or not use a decoy, and the aircraft strategies are to either pursue or not pursue. The game ends when the actual launch vehicle moves: if the aircraft is overhead, the vehicle is destroyed; otherwise, the missile launch is successful. The study shows that the optimal strategy for the players depend on whether or not the aircraft's past decisions are observable. If they are observable, decoys provide a significant advantage to the ground vehicle. If they are not observable, the decoys offer no advantage.

Wagner and Arkin [86] explored the possibility of deception in robotics and showed that a better model of the opponent improved the effectiveness of deception. The authors perform an experiment where two robots play hide and seek. Each robot models its environment as a two-player, zero-sum game. The seeker earns a positive payout for locating the hider. The hiding robot can send deceptive signals to the seeker by knocking down dry-erase markers placed along its path. These signals are meant to manipulate the seeker's perception of the true game's payouts. The seeking robot is equipped with a simulated suite of sensors, e.g. infrared, auditory, or visual. The seeker's ability to observe signals from the hider is based upon the number of sensors in its suite—the more sensors, the better the seeker was able to observe the environment. By searching an area with a toppled marker, the seeker is more likely to obtain a higher payout than by searching an area where all the markers remain upright. Some information about the seeker's sensors is revealed to the hider before each experimental trial, e.g. that the seeker has an infrared sensor or that it does not have visual sensors. The results show the hider deceives more effectively when given additional information about the seeker's sensors: the rate at which deception succeeded increases from 78 percent when the hider had no information about the seeker's sensor suite to 95 percent when it has full knowledge of the seeker's sensor suite.

2.6 Responding to Deception: Mitigation and Counterdeception

Much of the research that focuses on performing deception also addresses ways of mitigating the effects of deception. Deception mitigation is applicable in cooperative mechanism design—an area of increasing importance in several emerging fields. For instance, mobile ad hoc networks (MANETs) rely on the ability of autonomous network nodes to operate in dynamic environments. An adversary can degrade network performance by introducing a malicious node. The malicious node could broadcast spurious signals to entice the honest MANET nodes to organize in a less efficient configuration. However, the MANET can protect against such malicious activities by using deception mitigation techniques. Although, game theory has been used to reduce the impact of malicious nodes in MANETs [59], additional work is needed to mitigate the risk of deception in MANETs and other emerging fields. Furthermore, no article explicitly addresses the use of counterdeception.

"Counterdeception strives to identify and exploit adversary attempts to mislead... [it] includes actions taken to force adversaries to reveal their. . . intentions and objectives" [77]. The difference between mitigation and counterdeception is this: mitigation seeks to avoid the worst-case outcome, counterdeception seeks to improve the current expected outcome by turning the deception against the deceiver.

The problem of counterdeception is difficult because it assumes the mark is able to (1) detect ongoing deception, (2) accurately model the opponent, and (3) formulate a response to the deceptive activity. Bennett and Vane [6] identifies three objectives that would need to be met for hypergames to be an effective planning tool for deception. These objective are relevant more generally to any tool used for either deception planning or counterdeception planning. The model must:

1. Allow planners to compare (counter)deception plans to deception-free plans
2. Allow planners to evaluate cost of potential failure
3. Focus attention on planner's belief contexts regarding the adversary's potential courses of action

Hypergame Theory, with its ability to model opponent perceptions to the k-th level, provide a straight-forward approach to such a problem. Furthermore, the Hypergame Normal Form (HNF) described in [78] is intentionally designed to accomplish these three objectives. Still, we, the authors, have seen no study that applies Game Theory, Hypergame Theory or the HNF to the problem of performing counterdeception.

The remainder of this section describes several studies that present optimal strategies of deception mitigation by a deception-aware mark.

Lee and Teo [50] present a simple signaling game that consists of two players and two boxes. Box 1 and box 2 are assigned values (i, j) from among a set S of possible values according to a collection of discrete probability functions, where p_{ij} is the probability of assigning the pair (i, j) to the boxes. The deceiver observes the box values, closes the boxes, and affixes a label (k, j) to the outside of the boxes. The mark then observes the label, selects a single box, and receives the payout according to the original value contained inside. The value of the game is denoted by v. Lee and Teo showed that if the following inequality holds as an equality, then the deceiver has completely distorted the information:

$$v \geq \max \left\{ \sum_{(i,j)\in S} i p_{ij}, \sum_{(i,j)\in S} j p_{ij} \right\}$$

On the other hand, if the inequality holds strictly, then the first player has not entirely distorted the information. In such a case, the second player should simply choose the box with the highest mean box contents. Thus, the authors demonstrate that a suboptimal deception can be overcome through careful analysis of the game's history.

In [52], Lisý*et al.* consider a game where an adversary camouflages targets in an area to degrade a mobile sensor team's ability to accurately perceive the targets' importance. Despite this obfuscation of priorities, the mobile sensor teams must arrange themselves to provide sufficient coverage of the targets. An algorithm is introduced that maximizes the worst-case performance of the sensor team, i.e. the algorithm guarantees a minimum level of performance despite deception. This approach is worthwhile in high-risk environments (e.g. security scenarios); however it is not generally useful to assume a worst-case scenario [50]. Future work should explore an alternative, hybrid solution that maximizes the mark's payoff unless and until the potential cost of deception reaches some predefined threshold. Once reached, the mark could enact a loss-prevention strategy with a MaxiMin approach. This hybrid solution might have amortized better performance when the penalty for being deceived is relatively low.

2.7 Conclusion and Future Work

Adversarial conflicts often give rise to deception. This chapter provides an overview of game theory as a model of deceptive conflicts. Section 2.3 describes the ways in which the practicality of deception is measured. First, the situation must warrant the use of deception. Next, the deceiver must have the capacity to deceive, and the conflict must be one of incomplete/imperfect information. Finally, section 2.3 describes several properties that were used to determine the desirability of a deception (i.e. stealth, surety and cost), and a list of deceptive mechanisms for games is presented.

Section 2.4 highlights a need for additional work focused on detecting deception in games. Section 2.5 introduces the various ways in which game theory is used to model deception in the literature. Section 2.6 surveys several studies that focused on deception mitigation. Still, much work is needed in this area. Furthermore, the discussion highlights the need for works that specifically model counterdeception, i.e. the ability to expose or exploit a deceptive adversary's actions. Two approaches are recommended for mitigating or countering deception that should be considered in future work: performing exploratory actions and value loss threshold-based strategies. Several other areas should be addressed in future studies, as well. Specifically, future efforts should consider the following:

- The benefit-risk trade-off inherent to every deception, especially against deception-aware opponents
- Deception as a costly activity either by parameterizing the payouts or using a multi-objective approach
- Situations where the deception uses multiple mechanisms simultaneously to create synergistic effects (e.g. payout manipulation combined with tactical deception)
- The effects of uncertainty on the part of the deceiver; many models assumed the deceiver had perfect or complete knowledge

III. A multi-objective optimization problem for environmental deception using game theory

3.1 Introduction

Adversarial conflicts often give rise to deception. Therefore, a model of conflict is needed to thoroughly study deception. Game theory is "the study of mathematical models of conflict" [56]. It is used in multi-agent conflicts to determine strategies that maximize an agent's expected utility when the utility function depends on the combined actions of all agents. Game theory has been used to study deception in many different ways, but there remain several gaps in the research, especially pertaining to the cost and risk of deception [21]. Deceptive cost describes the amount of effort required to cause an opponent to act according to the projected view of the conflict. Risk describes the potential loss when playing an off-equilibrium deceptive strategy against a deception-aware opponent; it depends on the accuracy of the opponent model used to counter the deception [13]. It is important to address the benefit of deception along with it's cost and risk because deception in real-world conflicts is commonly realized as a trade-off between these three objectives. However, modeling cost and risk is difficult: It is hard to formulate the benefit, cost, and risk in a single utility function, and parameterized utility functions are not a general solution since these objectives are not always comparable. It, therefore, comes as no surprise that most of the game theoretic (GT) literature considers only cost-free deceptions. Furthermore, a recent survey of GT models of deception [21] indicates that risk is entirely absent from the literature. This is surprising because risk assessment is a critical factor in deception planning doctrine [77].

In light of these gaps, this chapter focuses on deception via payout manipulation in normal form games. Payout manipulation, also called environmental deception, is achieved when the deceiver successfully causes the target player (called the mark) to misperceive one or more of the game's utility functions. This chapter presents a GT model of environmental deception and formulates the task of designing efficient deceptive payout matrices as a multiobjective optimization problem—the deception design problem (DDP). A solution to the DDP is a payout matrix which, if presented to the mark as the true state of the conflict, is Pareto-efficient according to benefit, cost, and risk for the deceiver. Several versions of the DDP objective functions are discussed to help researchers articulate the deceiver's goals. A case study demonstrates the DDP using a 7×7 normal-form game. The game's payouts are derived from an air-to-air combat simulator [63].

The remainder of this chapter is organized as follows. Section 3.2 introduces the Environmental Deception Game used in the DDP. Section 3.3 defines the DDP as a multiobjective optimization problem. Section 3.4 describes the methodology for solving the DDP and presents an example based on the classic Prisoner's Dilemma. Section 3.5 presents a case study using a 7×7 normal form game. Related works are described in Section 3.6. Concluding remarks and areas for future work are given in Section 3.7.

3.2 The Environmental Deception Game

Environmental deception is a mechanism that causes the mark to misperceive the conflicts' payouts. Since the payouts express player preferences, misperception of the payouts can cause the mark to select a suboptimal strategy. The deceiver can take advantage of the mark's misperception by anticipating their response and responding accordingly. The Environmental Deception Game (EDG) provides a model for this kind of deception. The EDG is a game defined by a tuple of payout matrices: the true payout matrix and the deceived payout matrix. The true payouts describe the actual state of the conflict; the deceived payouts are generated by the deceiver to mislead the mark. The game assumes, at the end of play, players observe their own received payout, but they do not know the true payouts received by their opponent.

The game with true payouts A and and deceived payouts B is denoted $G = \langle A, B \rangle$. Let a_{ij}^{p} denote the payout for player p at the intersection of row i and column j in the $m \times n$ payout matrix A. Given strategy profiles s and t for the row and column player respectively, the expected utility for the row player is

$$E_{row}(G) = \sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij}^{row} s_i t_j \tag{Row Expected Utility}$$

and the expected utility for the column player is

$$E_{col}(G) = \sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij}^{col} s_i t_j \tag{Column Expected Utility}$$

For consistency, the row player is the deceiver and the column player is the mark. The deceiver plays according to the true payouts in A while the mark plays according to the deceived payouts in B. The mark believes both players are playing the same game, i.e. according to the Nash equilibrium of the deceived payouts. However, the deceiver plays a different game: The deceiver anticipates the mark's strategy according to the deceived payouts and computes the best response according to the true payouts. In this way, the deceiver takes advantage of the mark's misperception of the conflict.

Denote by σ_X^p the strategy profile for player p according to the Nash equilibrium of the matrix X. Since the mark plays according to the payout matrix B, the mark's strategy profile is denoted, σ_B^{mark}. The deceiver's best response to σ_B^{mark} is called the *deceptive strategy* and denoted σ_D . The deceptive strategy $\sigma_D = \mathbf{X} = (x_1, x_2, \ldots, x_m)$

is be computed by solving the following linear program:

$$\begin{aligned}\mathbf{X} \; = \; & \text{maximize}_x \;\; \mathbf{C}x \\ & \text{subject to} \;\; \sum_{i=1}^{m} x_i = 1, \\ & \qquad\qquad\qquad 0 \leq x_i \leq 1, \quad i = 1, \ldots, m\end{aligned}$$

where $\mathbf{C} = \{c_1, c_2, \ldots c_m\}$ and $c_i = \sum_{j=1}^{n} \left((\sigma_B^{mark})_j \cdot a_{ij}^{deceiver} \right)$ for $i = 1, \ldots, m$, with $(\sigma_B^{mark})_j$ being the j^{th} element of mark's strategy profile at the Nash equilibrium according to the deceived payouts B, and $a_{ij}^{deceiver}$ being the deceiver's payout in the ith row and jth column of the true payout matrix A.

Likewise, the mark's best response to σ_D is called the *counterdeception strategy* and denoted σ_{CD}. The mark's counterdeception strategy $\sigma_{CD} = \mathbf{Y} = (y_1, y_2, \ldots, y_n)$ is computed by solving the following linear program:

$$\begin{aligned}\mathbf{Y} \; = \; & \text{maximize}_y \;\; \mathbf{D}y \\ & \text{subject to} \;\; \sum_{j=1}^{n} y_j = 1, \\ & \qquad\qquad\qquad 0 \leq y_j \leq 1, \quad j = 1, \ldots, n\end{aligned}$$

where $\mathbf{D} = \{d_1, d_2, \ldots d_n\}$ and $d_j = \sum_{i=1}^{m} \left((\sigma_D)_i \cdot a_{ij}^{mark} \right)$ for $j = 1, \ldots, n$, with σ_D being the deceiver's's best response to the mark's strategy profile at the Nash equilibrium according to the deceived payouts in B, and a_{ij}^{mark} being the mark's payout in the ith row and jth column of the true payout matrix A.

3.3 The Deception Design Problem

The DDP is intended as a means of analyzing the benefit-cost-risk trade-off of environmental deception in multi-agent conflicts. As such, it is designed to answer the following question: Suppose one agent could successfully deceive another through environmental deception; how should such a capability be used to achieve an efficient trade-off between benefit, cost, and risk? Military utility analysts often consider questions of this kind in the earliest phases of a capability's acquisition life-cycle. They assume the capability exists and then try to predict its impact on operations. In the same way, this chapter is interested not in engineering a capability that performs environmental deception. Instead, it frames a discussion on efficient employment of such a capability and provides a means of predicting its impact. The Environmental Deception Game predicts the outcome of a conflict with one-sided misperception of the payouts. The DDP asks: If the row player (deceiver) could influence those misperceptions, how should they be changed?

The DDP makes two principal assumptions with regards to uncertainty. First, the deceiver's environmental deception is assumed to be successful, i.e. the mark plays according to the Nash equilibrium of the deceived payouts. Second, the deceiver is assumed to know the true payouts of the game. Relaxation of these assumptions is discussed in the future work section.

With this in mind, the DDP for two players is defined as follows. Given a payout matrix for a two-player game, A, compute a payout matrix, B for the environmental deception game $G = \langle A, B \rangle$ that maximizes benefit (f_B) and minimizes cost (f_C) and risk (f_R). The functions f_B, f_C, f_R are defined below. For notational convenience, the solution where $G = \langle A, A \rangle$, is called the null solution and is denoted $s_\emptyset$.

Benefit.

Benefit describes the increase in expected utility for using B as the environmental deception. Benefit is defined as

$$f_B(G) = E_{row}(G) - E_{row}(A), \qquad \text{(Benefit)}$$

where $E_{row}(G)$ is the deceiver's expected value in G and $E_{row}(A)$ is the expected value of the original game. Benefit compares the value of playing the environmental deception game G against playing the non-deceptive game according to A. The value of the null solution $s_{\emptyset}$ is $f_B(s_{\emptyset}) = 0$, since $G = \langle A, A \rangle$ implies $E_{row}(G) = E_{row}(A)$. On the other hand, when $A \neq B$, $f_B(G) \geq 0$, since misperception by the mark can cause him to deviate from the Nash equilibrium of the true game. The mark's unilateral deviation provides an opportunity for the deceiver to increase his expected utility over that of the Nash equilibrium by playing a best response. However, the deception provides zero benefit if the misperception of the mark does not induce changes in his equilibrium strategy.

Cost.

The cost function describes the amount of effort that must be expended to cause the mark to believe the deception. It is assumed that the cost of the deception varies according to the difference between the true payout matrix and the deceived payout matrix, i.e. the distance between the two. Thus, it is sensible to use a metric function to define cost:

$$f_C(G) = \mathbf{dist}(A, B), \tag{Cost}$$

where $\mathbf{dist}(A, B) : \mathcal{M} \times \mathcal{M} \rightarrow [0, \infty)$ is a metric, $\mathcal{M}$ is the set of all real-valued $m \times n$ payout matrices, and $[0, \infty)$ is the set of non-negative real numbers. Table 1 presents several cost metrics and domains in which they might be useful. The metrics in Table 1 are defined so that the cost of deceiving any particular payout is treated the same as those in any other payout. In some circumstances, it may be appropriate to measure cost differently based on the particular payout being changed, i.e. when the cost varies based on both the change to the true payouts *and* the outcome payout being modified. This approach is useful if it were possible to deceive on some payouts

Table 1. Example cost metrics

Cost Metric	Problem Domain
0	Cost-free deceptions
$\sum_{ijp}\left\lvert a_{ij}^{p}-b_{ij}^{p}\right\rvert$	Cost of deception is constant regardless of payout value
$\sum_{ijp}\sqrt{\left(a_{ij}^{p}-b_{ij}^{p}\right)^{2}}$	Small payout manipulations cost proportionately less than large payout manipulations
$\sum_{ijp}\left\lvert\int_{a_{ij}^{p}}^{b_{ij}^{p}}\frac{2}{1-\lvert 2x-1\rvert}dx\right\rvert$	Payouts are constrained to a range (e.g. [0,1]), and cost grows asymptotically near the limits of a payout's range

cost-free while other payouts in the same game are not cost-free.

Risk.

Deception sometimes involves an element of surprise, but often the best deception is the tacit one—the one that remains undiscovered even after many encounters. At what point does such a tacit deception pose a risk to the deceiver? It poses a risk if the mark discovers the deception and can formulate a counterdeception strategy before the deceiver changes his own strategy. Considering risk in a game where the deception is assumed to be successful may appear at first glance to be nonsense, but there is a potential for risk if the deception is revealed through the gameplay itself. If the mark is surprised by the outcome, he may question whether or not he perceived the conflict accurately. Surprise could spur an update to his belief state that uncovers the deception. If he discovers the true payouts unbeknownst to the deceiver, he could formulate the counterdeception strategy (σ_{CD}) that defeats the deceiver in subsequent encounters.

Risk assessments usually involve calculating two critical components: the magnitude of a potential consequence and the likelihood of its occurring [80]. For DDP, the magnitude of the consequence depends on the deceiver's level of exposure when playing an off-equilibrium strategy. Consequence is measured as the deceiver's value loss in the worst-case scenario, i.e. when the deceiver assumes the mark is successfully deceived, but the mark counters the deception perfectly. Thus, consequence is expressed as:

$$R_C = E_{row}(A) - \left(\sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij}^{deceiver} (\sigma_D)_i (\sigma_{CD})_j \right) \tag{Consequence}$$

where $(\sigma_{CD})_j$ is the j^{th} element of the mark's counterdeception strategy and $(\sigma_D)_i$ is the i^{th} element of the deceiver's deception strategy.

DDP assumes the deception is initially successful and that the deception is revealed if the payouts do not match the mark's perception of the game. The likelihood component of risk is based on the probability that the game's actual outcome and the mark's perception of the conflict are inconsistent, i.e. that the deception is revealed when the game is played. For the deception to be revealed, two conditions must be satisfied. First, the mark must receive a payout. Given strategy profile $s = (s_1, \ldots, s_m)$ for row and strategy profile $t = (t_1, \ldots, t_n)$ for column, the probability of receiving outcome a_{ij} equals $s_i \cdot t_j$ and is denoted $p_{ij}(s, t)$. The second condition is that the payout received by the mark must contradict his perception of the outcome's payout. The discrete indicator function $\mathbf{I} : \mathbb{R} \times \mathbb{R} \to \{0, 1\}$ is used to indicate whether or not the mark's perception of the game agrees with the payout he receives. If $x = y$, then $I(x, y) = 0$. Otherwise, $I(x, y) = 1$. Thus, likelihood is defined by the following equation:

$$R_L = \sum_{ij} \mathbf{I} \left(a_{ij}^{mark}, b_{ij}^{mark} \right) \cdot p_{ij}(\sigma_D, \sigma_B^{mark}) \tag{Likelihood}$$

Given an environmental deception game G with likelihood of detection $R_L(G)$ and consequence $R_C(G)$, risk is expressed as the product:

$$f_R(G) = R_L(G) \cdot R_C(G) \tag{Risk}$$

To show how risk is affected by consequence, consider the situation where the deceiver changes all the payouts i.e. when $\forall i, j, I(a_{ij}, b_{ij}) = 1$. In this case, regardless of the strategies adopted by the players, the likelihood of discovery equals

$$\sum_{ij} 1 \cdot p_{ij}(s,t) = 1 \cdot \sum_{ij} p_{ij}(s,t) = 1$$

Thus, the deception is certain to be discovered. If the deception strategy σ_D is off-equilibrium, then the deceiver is exposed to the counterdeception strategy σ_{CD} and $R_C > 0$. In this case, $f_R = R_L \cdot R_C = 1 \cdot R_C > 0$. On the other hand, if the deception strategy σ_D is on-equilibrium, the deceiver does not change his strategy from the original game. Thus, the consequence is zero. As a result, $f_R = R_L \cdot R_C = 1 \cdot 0 = 0$.

To show how risk is affected by likelihood, consider the situation where the deceptive strategy σ_D is off-equilibrium. Then the consequence is non-zero. If the deceiver achieves the deception by changing only those payouts outside the deception game's equilibrium, then the players will only receive payouts that correspond to the original game's payout values. In this case, the likelihood of discovery is zero. Thus, $f_R = R_L \cdot R_C = 0 \cdot R_C = 0$. On the other hand, if the payouts in the deceived game's equilibrium have been changed, then $R_L > 0$ and $f_R = R_L \cdot R_C > 0$.

A special case arises in when the true payouts contain one or more correlated equilibrium. A correlated equilibrium is a generalization of the Nash equilibrium whereby no player has incentive to deviate from the equilibrium assuming the other player selects a strategy based on the same public signal—in this case, the deceptive payouts. Thus, if the true payouts have a correlated equilibrium and σ_D plays the correlated equilibrium, then the counterdeception strategy is to play the correlated equilibrium, as well. By playing a deception that follows a favorable correlated equilibrium, the deceiver is no worse-off for being discovered since (assuming rationality of the mark) the deceptive payouts force the mark's strategy choice. In this way, it is possible for R_C to be less than zero. Thus, for games where the only equilibrium is the Nash equilibrium, R_C has a straightforward interpretation: it is the expected value loss when the mark plays the counterdeception strategy. In games with correlated equilibria other than Nash, a negative value for R_C must be interpreted as forcing cooperation between the players.

To prevent ambiguity in the results, each solution lists the risk objective function value as well as the likelihood R_L and consequence R_C.

3.4 Methodology

This section outlines the process used to perform the case study in Section 3.5. A small demonstration using the prisoner's dilemma provides a concrete example.

Process.

Solving an instance of the DDP involves computing the best-known a set of non-dominated solutions, called PF_{known}. The Speed-constrained Multiobjective Particle Swarm Optimization (SMPSO) algorithm is used [58]. The SMPSO algorithm is discussed in Section 3.6.

The authors implemented the DDP and its associated objective functions using the MOEA Framework [34]—a Java library for multiobjective evolutionary algorithms. The equilibrium for each game is computed using the Enumeration of Extreme Equilibrium algorithm [66] and the linear programs are implemented using the JOptimizer version 3.4.0—a Java library for solving convex optimization problems. The true and

Table 2. SMPSO Settings

Setting	Value
Number of seeds	1000
Population size	100
Generations	250
Polynomial Mutation rate	1/100
Polynomial Mutation distribution index	20

deceived payouts are stored as $2 \times n \times m$ arrays. To account for finite representation of the solutions, the deceived payouts are rounded to the nearest $10^{-5\text{th}}$ decimal place before evaluating the objective functions.

SMPSO is executed 1,000 times using a unique random seed for each run. The settings for SMPSO are listed in Table 2. Each run produces a set of nondominated solutions including their associated fitness values, the deceiver's deceptive strategy σ_D, and the mark's optimal counteredeception strategy σ_{CD}. All 1,000 solution sets are inserted into an intermediate set, which is then reduced to PF_{known} by removing all dominated solutions. Furthermore, any solution dominated by the null solution, $s_\emptyset$, is removed. The solutions in PF_{known} are then decoded back into payout matrix form and analyzed.

Example: Prisoner's Dilemma.

An instance of the prisoner's dilemma illustrates the process. The prisoner's dilemma is a classic problem presented in introductory game theory courses; the payouts are shown in Table 3.

It is often introduced as a conflict between two prisoners with regards to a plea bargain. The prisoners agreed that if caught, they would not confess the crime. However, they have been caught and separated so that neither knows what the other will choose. Each must decide to either cooperate with their fellow prisoner or to

Table 3. The Prisoner's Dilemma

	Cooperate	Defect
Cooperate	(-1,-1)	(-3,0)
Defect	(0,-3)	(-2,-2)

defect, i.e. to confess their crime to the police. If both cooperate, they face only a one-year sentence for a lesser crime. If they both defect, their confessions earn them a two years sentence. But if only one prisoner defects, his confession will be used against the other: The defecting prisoner is released immediately, and his fellow prisoner will face the maximum sentence of three-years. The game exhibits a pure strategy Nash equilibrium wherein both prisoners confess and receive two-year sentences. This equilibrium is underlined in Table 3.

In this example, the cost of deception is assumed to be constant regardless of the payout value. Therefore, the L^1-norm is used as the cost metric, i.e.

$$f_C(G) = \sum_{ijp} \left| a_{ij}^{(p)} - b_{ij}^{(p)} \right| .$$

After executing SMPSO and removing all dominated solutions from the solutions sets, PF_{known}, contains only a single solution s_1. The deceived payouts B defined by s_1 are shown in Table 4; differences from the true payouts are shown in bold. s_1 exhibits a pure strategy Nash equilibrium at (Defect, Cooperate), underlined in Table 4. The deceiver's best response is to Defect, i.e. $\sigma_D = (0, 1.0)$. The expected utility of the original game matrix A is -2 since both players receive a two-year sentence. In the environmental deception game $G = \langle A, B \rangle$, the deceiver's expected utility is 0 and the mark's expected utility is -3. Since the deceiver improves his expected utility by two years, the benefit of deception is $f_B(G) = 2$.

Table 4. DDP Example Solutions. Deceived payout matrix B for solution s_1 to the Prisoner's Dilemma. Deviations from the true payouts are shown in bold. The equilibrium of the EDG is underlined.

	Cooperate	Defect
Cooperate	(-1,**0**)	(-3,0)
Defect	(0,**-1.999**)	(**-1**,-2)

Since the L^1-norm is used as the cost metric, the total cost of s_1 is:

$$f_C(G) = \sum_{ijp} \left| a_{ij}^{(p)} - b_{ij}^{(p)} \right| = (1 + 1.001 + 1) = 3.001$$

Given the mark is deceived, the mark plays Cooperate while the deceiver plays Defect. As a result, the mark expects a payout of -1.999, but receives -3 instead. This result is certain since the equilibrium of G is a pure strategy. Therefore, the likelihood component of risk $R_L = 1.0$. On the other hand, the mark's counterdeception strategy is to defect, i.e. $\sigma_{CD} = (0, 1.0)$. If the mark plays σ_{CD}, then the game returns to the original equilibrium, (Defect, Defect). Since the deceiver plays exactly as he would in the deception-free game, the consequence component of risk $R_C = 0$. Thus, the risk for performing this deception is $f_R(s_1) = 1.0 \cdot 0 = 0$. As an interesting result, this means that in a single-shot prisoner's dilemma, there is no reason not to deceive your opponent.

3.5 Case Study

This section presents a case study to demonstrate the DDP in a more complex game. Specifically, the case study uses a 7×7 normal form game based on the output

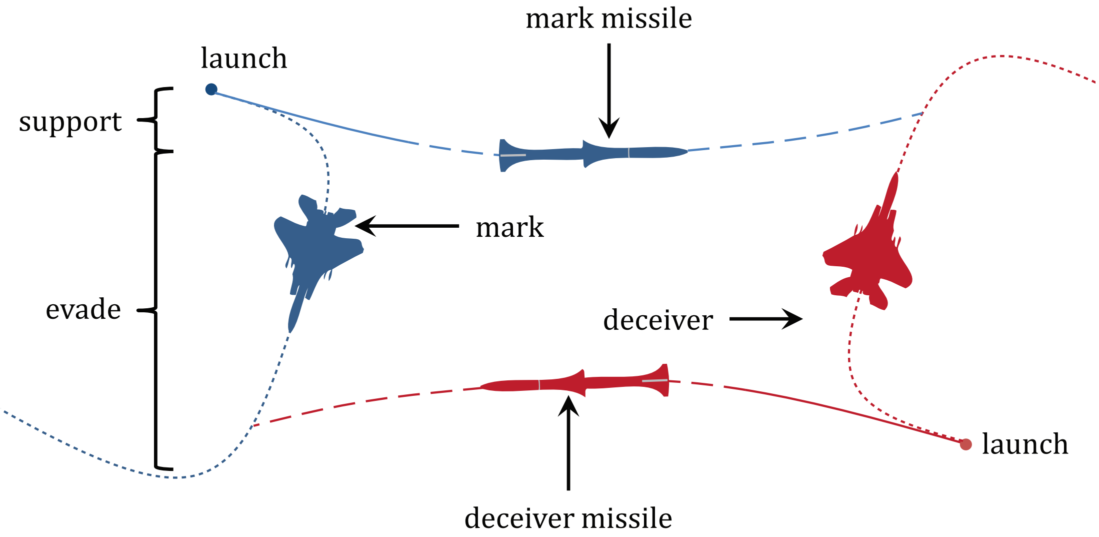

Figure 7. Missile Support Scenario. The aircraft must decide how long they will support their missile before evading the oncoming missile.

of an air-to-air combat simulator, the Missile Support Time (MST) game [63], shown in Figure 7. The conflict involves two identical aircraft (AC) approaching each other at the edge of their engagement range. They each fire a missile at their opponent. The pilots can increase their probability of kill (P_K) by supporting their missile with updates from their AC's guidance system. The pilots can enhance their chances of survival by performing evasive maneuvers. A pilot can either evade or support, but not both. To evade, the pilot must break his lock on the opponent and cease supporting his missile. So, the pilots must decide how long they will support before they evade. The longer they support, the more likely they are to kill their target, but this exposes them to a higher probability of being shot down themselves. Conversely, if they evade too soon, they are more likely to survive the engagement, but their missile will not likely hit their opponent.

This case study investigates the possibility of using deception to cause one pilot to either support too long, or evade too soon. The payouts for each pilot are defined by a regression model fitted to the output of a discrete-event air combat simulator [63].

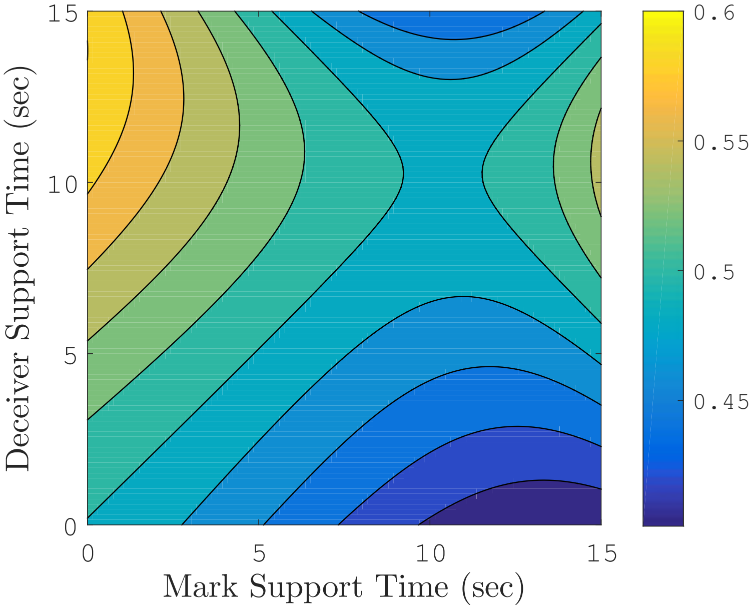

Figure 8. Deceiver Payouts for Missile Support Time game. The payouts for the mark and deceiver are nearly symmetric. The game exhibits a pure strategy Nash equilibrium wherein both players support for ten seconds.

The payouts for the players are nearly symmetric, and the payouts for the deceiver are shown in Figure 8.

The payouts are determined by a regression model of the form

$$p(x, y; \beta) = \frac{\exp(q(x, y; \beta))}{1 + \exp(q(x, y; \beta))}$$

where x is the support time of the deceiver, y is the support time of the mark, and $q(x, y; \beta)$ is a quadratic function of decision variables x and y, i.e.

$$q(x, y; \beta) = \beta_0 + \beta_1 x + \beta_2 y + \beta_3 x^2 + \beta_4 y^2 + \beta_5 xy.$$

The resulting P_K values are based on the β parameter vectors for the deceiver (β_D) and mark (β_M) presented in Table 5.

Table 5. Parameter vectors for the deceiver (β_D) and mark (β_M)

variable	parameter	$\boldsymbol{\beta_D}$ Deceiver	$\boldsymbol{\beta_M}$ Mark
constant	β_0	-3.440	-3.529
x	β_1	0.289	-0.013
y	β_2	-0.131	0.300
x^2	β_3	-0.009	0.011
y^2	β_4	0.012	-0.009
xy	β_5	0.003	0.003

Thus, given that the deceiver supports for x seconds and the mark supports for y seconds, the deceiver's P_K is $p(x, y; \beta_\mathbf{D})$ and the mark's P_K is $p(x, y; \beta_\mathbf{M})$. Since the probability of survival equals $(1 - P_K)$, the probability of survival for the deceiver is $(1 - p(x, y; \beta_\mathbf{M}))$, and the probability of survival for the mark is $(1 - p(x, y; \beta_\mathbf{D}))$. A weighted sum defines each player's payout as a trade-off between achieving a kill and surviving the engagement, i.e.

$$u_{deceiver}(x, y) = \omega_D p(x, y; \beta_\mathbf{D}) + (1 - \omega_D)(1 - p(x, y; \beta_\mathbf{M}))$$

and

$$u_{mark}(x, y) = \omega_M p(x, y; \beta_\mathbf{M}) + (1 - \omega_M)(1 - p(x, y; \beta_\mathbf{D}))$$

where $0 \leq \omega_M$ and $0 \leq \omega_D$. For this case study, ω_D and ω_M are set to 0.5.

The strategies x and y are discretized so the game can be represented in normal form. The resulting support times range from 0 seconds to 15 seconds at 2.5 second intervals, i.e. the pilots can support for either 0s, 2.5s, 5s, and so on. The maximum support time of 15 seconds is based on the maximum flight time of the simulated missiles. Under these conditions, the game exhibits a pure strategy Nash equilibrium when each player supports for 10 seconds before evading. As in the prisoner's dilemma example, this case study uses the L^1-norm as the cost function.

Table 6. Solutions to the Missile Support Game.

Index	Fake Game Strategy Profile of Mark (σ_B^{mark})							σ_D	σ_{CD}	f_B	f_C	f_R	R_C	R_L
	0s	**2.5s**	**5s**	**7.5s**	**10s**	**12.5s**	**15s**							
s_1	1.00	-	-	-	-	-	-	15s	10s	**0.095**	3.263	**-**	0.048	-
s_2	1.00	-	-	-	-	-	-	15s	10s	**0.095**	3.159	0.048	0.048	1.000
s_3	0.70	-	0.30	-	-	-	-	12.5s	10s	0.076	3.152	0.008	0.008	1.000
s_4	-	-	-	-	-	-	1.00	10s	10s	0.056	3.151	**-**	-	1.000
s_5	-	-	0.67	-	-	-	0.33	10s	10s	0.039	3.148	**-**	-	0.328
s_6	-	-	0.67	-	-	0.01	0.32	10s	10s	0.039	3.133	**-**	-	1.000
s_7	-	-	0.64	0.30	-	0.06	-	10s	10s	0.024	**3.103**	**-**	-	1.000

SMPSO found 71,415 solutions total during the 1,000 runs. After the dominated solutions are eliminated, PF_{known} contained seven non-dominated solutions. The solutions are enumerated in Table 6, sorted by benefit f_B. The first column lists the index of each solution. The next seven columns show the mark's equilibrium strategy based on the deceived payouts of the solution. Thus, the deceived payouts for s_1 exhibit a pure strategy Nash equilibrium where the mark supports for 0 seconds. The deception strategy for the deceiver is listed in the σ_D column, and the counterdeception strategy for the mark is listed in column σ_{CD}. For s_1, the best response to the mark is for the deceiver to support for 15 seconds; the best counterdeception is for the mark to support for 10 seconds. Equilibrium strategies throughout the table are rounded to the nearest 100$^{\text{th}}$ decimal place for presentation purposes. The objective function values for each solution are listed in the adjacent three columns with the most desirable values for each function shown in bold. Since risk is the product of consequence R_C and likelihood R_L, the final two columns display the consequence and likelihood components for each solution.

The most beneficial solution, s_1, provided an increase in the deceiver's expected utility of 0.095. The payout manipulations cause the mark to support for 0 seconds (i.e. the fire and forget strategy). The deceiver can take advantage of this and support his missile for a full 15 seconds. So, the EDG results in the outcome at (15s, 0s).

s_1 is the highest cost solution: the absolute deviations from the true payouts summed to 3.263. s_1 modified 97 of the 98 ($= 7 \times 7 \times 2$) payout values in the game. The only payout that remained untouched was the payout for the mark at (15s, 0s), i.e. the expected outcome when the mark plays σ_B^{mark} and the deceiver plays σ_D. The solution increased the values of 58 payouts and decreased the values of 39 payouts. The solution applies a large portion of the cost (+0.149) to the payout in outcome (15s, 7.5s) for the mark. Since s_1 leaves the mark's payouts in (15s, 0s) untouched, it achieves zero likelihood of detection (R_L). However, if by some alternative means, the mark discovers the deception, the counterdeception strategy is to support for 10s; the consequence (R_C) in this case is 0.048 for the deceiver. Recall $f_R = R_L \cdot R_C$. Thus, $f_R(s_1) = 0 \cdot 0.05 = 0$.

The risk to the deceiver was zero in five of the seven cases. Four of the zero-risk strategies (s_4, s_5, s_6, s_7) achieved zero risk because the deception strategy σ_D is to support for 10 seconds. Since supporting for 10 seconds is the equilibrium strategy of the original game, the deceiver faces no consequence for playing these strategies. s_1 is the only other zero-risk solution. No solution had both zero consequence and zero likelihood of detection. Still, the observed risk values are lower than anticipated; it was surprising that so many zero- and low-risk solutions could be located in a game with 98 ($= 7 \times 7 \times 2$) free variables.

In all cases, σ_D implied the deceiver should support for at least 10s before evading, i.e. at least as long as in the deception-free game. Surprisingly, the mark's counterdeception strategy, σ_{CD}, is to revert to the true game's equilibrium. However, that σ_{CD} reverts back to the original equilibrium is not true for every possible deceptive strategy—only the non-dominated solutions found in this case study. If the deceiver supported for less than 10 seconds, the mark responds best with a longer, off-equilibrium support time. For instance, if the deceiver supports for 0 seconds, the counterdeception strategy is $\sigma_{CD} = 15s$.

The tendency to support longer when deceiving is influenced by the pilot preferences weights, ω_D and ω_M. However, the tendency to cause the mark to support for a shorter amount of time applies generally. If the deceiving pilot favors kills over survival, then the deception improves utility when the pilot can support longer. This is possible if the mark evades sooner. On the other hand, if the deceiving pilot favors survival over kills, then the deception improves utility when the mark supports for less time. Thus, in both cases, the tendency is to deceive the mark into evading sooner.

The lowest-cost solution was s_7 with $f_C(s_7) = 3.103$; however, this solution also had the lowest benefit, as well. One interesting observation is that there is no incentive for the algorithm to reduce the number of altered payouts, only the total magnitude of the changes. It's not surprising that this algorithm produces solutions where very few payout values remained unchanged. Specifically, each solution changed at least 94 payout values. Plotting the number of changed payouts according to the magnitude of change shows that half of the payouts are modified by less than 0.02 utility points. This relationship is shown in Figure 9. The light-grey lines in Figure 9 are the traces of each solution, and the black line is the average for all seven solutions. The relationship between the magnitude of manipulation and the number of changed payouts indicates that the solutions produced by SMPSO may be noisy.

3.6 Related Works

The benefit and cost objective functions resemble those found in [64], which presents an algorithm for computing desirable costly environmental deceptions in 3×3 normal form games. The risk objective function was based on the principles of risk assessment [80]. The effects of misrepresenting preferences was studied in two-player, 2×2 games [14] and in simple three-player voting games [16]. That the

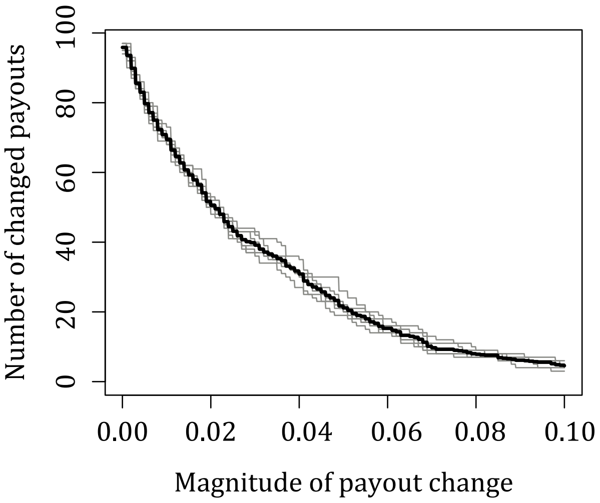

Figure 9. Number of changed payouts in DDP. The number of changed payouts are shown according to the magnitude of the payout's change. The light-grey lines are the traces for solutions $s_1, s_2, \ldots, s_7$. The black line shows the average of all ten solutions.

deception and counterdeception strategies are pure strategies coincide well with the results of [79] which shows that a pure strategy is often the best response to an opponent who misperceives the conflict. The environmental deception game is partially inspired by Hypergame theory [7]. It differs in that one player is assumed to know the true state of the conflict; Hypergame theory does not make this assumption in general.

The algorithm used to compute PF_{known}, SMPSO, is based on Particle Swarm Optimization (PSO), which is inspired by the flocking behavior of birds [46]. Although the no free lunch theorem implies that no one single optimization algorithm works best on all problems [89], PSOs have been proven to be effective tools for solving both continuous nonlinear and discrete binary optimization problems [24, 46–48]. In PSO algorithms, each solution is called a *particle*, and the population of solutions is called the *swarm*. Each particle i has a position $\vec{x_i}$ and velocity $\vec{v_i}$. The PSO updates the position at each generation t using the formula

$$\vec{x_i}(t) = \vec{x_i}(t-1) + \vec{v_i}(t)$$

where $\vec{v_i}(t)$ is given by

$$\vec{v_i} = w \cdot \vec{v_i}(t-1) + C_1 \cdot r_1 \cdot (\vec{x_{p_i}} - \vec{x_i}) + C_2 \cdot r_2 \cdot (\vec{x_{g_i}} - \vec{x_i}),$$

and $\vec{x_{p_i}}$ is the best solution according to the ith particle, $\vec{x_{g_i}}$ is the best solution known to the swarm, w is the inertia weight, and C_1 and C_2 dictate the influence of inertia and swarm particle attraction on the velocity of x_i.

SMPSO [58] constrains the velocity of each variable (in each particle) to prevent the particle velocities from "exploding" [22]. Constraining the velocity allows the algorithm to perform well even on difficult multiobjective problems (e.g. ZDT4 [93]). The pseudocode for SMPSO is given in Algorithm 1.

Algorithm 1 SMPSO pseudocode [58]

```
initializeSwarm()
initilizeLeadersArchive()
generation = 0
while generation < maxGenerations do
    computeSpeed()
    updatePosition()
    polynomialMutation()
    fitnessEvaluation()
    updateLeadersArchive()
    updateParticlesMemory()
    generation++
end while
return LeadersArchive()
```

3.7 Conclusion and Future Work

A recent survey of game theoretic models of deception [21] highlights a gap in GT literature regarding deception in the trade space of between benefit, cost, and risk. This chapter presents a novel approach to fill this gap. Section 3.2 introduces a GT model for environmental deception in situations where an opponent's perception of the conflict's payouts can be altered by the deceiver. The DDP is introduced in Sectoion 3.3 as a multiobjective optimization problem to design efficient environmental deceptions regarding benefit, cost, and risk. Benefit is measured as the value difference between the non-deceptive and deceptive games. Cost is computed using a metric to measure the distance between the original payouts and the deceived payouts. Risk is the product of the likelihood of discovery and consequence of discovery [80]. When the true game payouts exhibit correlated equilibrium other than the Nash equilibrium, the measure of risk is useful, but it can be harder to interpret. If risk likelihood and consequence are treated as separate objective functions, this difficulty is eliminated.

Section 3.4 outlines the process used to perform the case study in Section 3.5 using the Missile Support Time (MST) game. The known Pareto front for the MST consists of seven solutions. The solutions are obtained by executing a multiobjective evolutionary algorithm, SMPSO, 1000 times. Each time, a population of 100 particles is evolved over 250 generations. Of the seven solutions, five exhibited zero risk indicating that SMPSO is able to effectively locate zero-risk solutions. Four solutions are zero-risk because there is no consequence for being discovered, i.e. the deception strategy coincides with the deceiver's strategy in the deception-free game. The fifth zero-risk solution has zero-valued risk likelihood component. Thus, the mark's expected outcome in the deceived payout matrix coincides exactly with the expected outcome in the true payout matrix. Concerning the solutions' costs, it is possible that some cost observed in the resulting solutions is an artifact of the algorithm, but this observation requires further investigation.

This chapter provides several opportunities for future exploration by relaxing the two principal assumptions made by the DDP. First, the DDP assumes that the deceiver's environmental deception is successful. An environmental deception is successful if two conditions are satisfied: 1.) the mark correctly perceives the deceived payouts provided by the deceiver, and 2.) the mark chooses to play the Nash equilibrium. This assumption could be relaxed to consider situations where one of these two conditions are not satisfied. If a relaxation allows the mark to incorrectly perceives the deceived payouts, researchers should characterize the degree to which the mark's perception deviate from the payouts provided by the deceiver. If a relaxation allows the mark to select a strategy other than the Nash equilibrium, it is necessary to define a reasonable alternative. The DDP also assumes that the deceiver accurately perceives the true payouts of the conflict. This assumption could be relaxed to explore the effects of misperception on the part of the deceiver. A more general form of the environmental deception game would need to be defined. A possible solution would be constructed of three matrices: One for the true payouts, one for the deceiver's perception of the conflict, and one for the deceived payouts presented to the mark. The strategies of the mark and deceiver could be selected as they are in the version of the environmental deception game; however, the expected utility for each player would be computed according to the true payouts (vice the deceiver's perception of the true payouts). In this way, the environmental deception game defined above becomes a special case of this more general version. If the DDP is used as a tool for post hoc analysis, this approach might be able to address the issue of misperception on the part of the deceiver. In an ad hoc situation, this approach could be used to perform sensitivity analysis on the solutions.

Finally, all solutions to the case study manipulate at least 94 out of 98 payouts, but most of the payouts change by less than 0.02. This case study used the L^1-norm as a cost metric. Researchers should evaluate the impact of the cost metric on the number of changed payouts, e.g. by comparing solutions with an L^1-norm cost to those that use L^p-norm cost with $0 < p < 1$. However, the cost metric should be carefully selected based on the problem domain. Therefore, future work should evaluate alternative techniques for reducing the total number of changed payouts, as well. One simple technique is to introduce an additional objective function that counts the number of changed payouts; however, the introduction of a fourth dimension in objective space will likely increase the number of non-dominated solutions considerably. Alternatively, future work can evaluate the performance of other optimization algorithms. Another technique is to adjust the cost (or benefit) objective function according to the number of changed payouts. For example, the cost can be scaled by multiplying cost by $\frac{1}{p-c}$, where p is the total number of payouts and c is the total number of changed payouts. As c approaches p, the cost increases, penalizing solutions with many changed payouts. One final approach, borrowed from the Ridge Regression shrinkage method [45], penalizes cost additively based on the magnitude of changed payouts. For instance, suppose the maximum allowable deviation is $\pm\delta_{\max}$. Then the cost can be penalized by adding the shrinkage penalty, $\lambda \sum_{ij}(1-\delta_{ij}/\delta_{\max})^2$, where δ_{ij} is the absolute deviation in row i and column j, and lambda is a tuning parameter. The shrinkage penalty is inversely proportional to the size of the deviation; it effectively shrinks the deviations toward zero. The tuning parameter λ controls the relative impact of the penalty on the resulting solutions. One final approach is to reduce solution noise mid- or post- execution. A greedy approach could reduce the smallest deviations first unless and until a change in the equilibrium is observed before moving to subsequently larger deviations.

IV. Conclusion

4.1 Conclusion and Path Forward

This research is presented as two scholarly articles. The first is a survey of deception in game theory and the second introduces a novel multiobjective approach to deception modeling. The concluding remarks and path forward for each article are discussed in turn.

Survey of Deception in Game Theory.

A survey of the various game theoretic models of deception is presented. It provides background on the information models and game classes used to study deception. The survey covers numerous models for detecting, performing, and responding to deception. However, the majority of the literature focuses on performing deception. Four topics require additional attention in future efforts.

The first topic relates to the benefit-risk trade-off of deception. The deceiver is exposed to risk when playing an off-equilibrium strategy. Little has been done to address the risk of deception, especially as a trade-off to benefit or value gain. Characterizing risk is an important area for future work because deception is often realized as a trade-off between benefit and risk.

Costly deception is the second topic that requires additional attention. Many of the deception models considered strictly cost-free deceptions. In some scenarios, cost-free deception is appropriate, but this is not true in general. One approach is to parametrize the utility functions based on the cost of deception. This common approach is rarely used in the deception literature. One reason for such little attention is that analyzing conflicts with parameterized utility functions can be difficult. Another reason is that cost and utility are not always comparable; in this case, it is inappropriate to parameterize the player utility functions based on the cost of deception. Thus, future work should consider deception from a multiobjective perspective, measuring cost and benefit separately. This approach is used in Chapter III to define the Deception Design Problem.

The survey also identifies an absence of models that use multiple deceptive mechanisms simultaneously. Recall, a deceptive mechanism describes the means by which the deception is accomplished from a game theoretic perspective. The survey identifies five deceptive mechanisms, namely environmental deception (payout manipulation), tactical deception, action deception, participant deception, and misrepresenting player types. When combined, the mechanisms may improve the effectiveness of the deception synergistically.

The literature rarely considers the effects of uncertainty on the deceiver. Instead, much of the literature uses an asymmetric model of information, i.e. researchers often assume the deceiver has perfect knowledge of the conflict. But in many real-world conflicts, this is seldom the case: there is often misperception on both sides of the deception. Future work should investigate models that allow misperception by the mark and also by the deceiver in order to measure the impact of uncertainty and misperception in a more realistic way.

Finally, the survey indicates that deception mitigation receives little attention in the literature. Deception mitigation seeks to minimize the effects of deception and has applicability in cooperative mechanism design. It is an area of increasing importance in several emerging fields. For instance, mobile ad hoc networks (MANETs) rely on the ability of autonomous network nodes to operate in dynamic environments. An adversary can degrade network performance by introducing a malicious node. The malicious node could broadcast spurious signals to entice the honest MANET nodes to organize in a less efficient configuration. However, the MANET can protect against such malicious activities by using deception mitigation techniques. Game theory has been used to reduce the impact of malicious nodes in MANETs [59], but additional work is needed to mitigate the risk of deception in MANETs and other emerging fields.

The Deception Design Problem.

The Deception Design Problem (DDP) is a novel approach to deception design, introduced as a multiobjective problem. The DDP addresses the benefit, cost, and risk trade-off observed in deceptive conflicts. The model uses a simple normal form game to define one-sided environmental deception. This game is called the Environmental Deception Game and resembles a two-player hypergame [7]. Definitions for the benefit, cost, and risk objective functions are presented and discussed. The DDP is implemented in Java and leverages the algorithms provided by the Multiobjective Evolutionary Algorithm (MOEA) framework [34].

A case study is used to demonstrate the viability of DDP to inform deception design. The case study uses a normal form version of the Missile Support Time (MST) game [63]. The Speed-constrained Multiobjective Particle Swarm Optimization (SMPSO) algorithm [58] produces solutions to the DDP that are Pareto efficient in terms of the benefit, cost and risk objective functions. After executing SMPSO 1,000 times, seven nondominated solutions were presented. The solutions describe the perceptions that are most desirable for the deceiver, in terms of their benefit, cost, and risk. Most of the solutions are zero-risk because they either have zero likelihood of detection or because the deceiver does not need to deviate from the equilibrium of the original game.

The DDP provides several opportunities for future research. First, the DDP assumes that the deceiver's environmental deception was successful, i.e. that the mark plays according to the Nash equilibrium of the deceived payouts. This assumption could be relaxed to consider situations where the deception's success is uncertain and the deceiver must operate in a more risk-averse manner. The second principal assumption was that the deceiver accurately perceived the true payouts of the conflict. This assumption could be relaxed to explore the effects of misperception on the part of the deceiver. A more general form of the environmental deception game would need to be defined. Finally, the case study solutions manipulate the majority of the payouts (at least 94), but most of the payouts change by less that 0.02. Future work should evaluate techniques for reducing the total number of changed payouts.

All solutions to the case study manipulate at least 94 out of 98 payouts, but most of the payouts change by less than 0.02. This case study used the L^1-norm as a cost metric. Researchers should evaluate the impact of the cost metric on the number of changed payouts, e.g. by comparing solutions with an L^1-norm cost to those that use L^p-norm cost with $0 < p < 1$. However, the cost metric should be carefully selected based on the problem domain. Therefore, future work should evaluate alternative techniques for reducing the total number of changed payouts, as well. One simple technique is to introduce an additional objective function that counts the number of changed payouts; however, the introduction of a fourth dimension in objective space will likely increase the number of non-dominated solutions considerably. Alternatively, future work can evaluate the performance of other optimization algorithms. Another technique is to adjust the cost (or benefit) objective function according to the number of changed payouts. For example, the cost can be scaled by multiplying cost by $\frac{1}{p-c}$, where p is the total number of payouts and c is the total number of changed payouts. As c approaches p, the cost increases, penalizing solutions with many changed payouts. One final approach, borrowed from the Ridge Regression shrinkage method [45], penalizes cost additively based on the magnitude of changed payouts. For instance, suppose the maximum allowable deviation is $\pm\delta_{\max}$. Then the cost can be penalized by adding the shrinkage penalty, $\lambda \sum_{ij}(1 - \delta_{ij}/\delta_{\max})^2$, where δ_{ij} is the absolute deviation in row i and column j, and lambda is a tuning parameter. The shrinkage penalty is inversely proportional to the size of the deviation; it effectively shrinks the deviations toward zero. The tuning parameter λ controls the relative impact of the penalty on the resulting solutions. One final approach is to reduce solution noise mid- or post- execution. A greedy approach could reduce the smallest deviations first unless and until a change in the equilibrium is observed before moving to subsequently larger deviations.

Bibliography

1. George Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84:488–500, 1970.

2. Hossein Arsham. Stability of essential strategy in two-person zero-sum games. In Ralph G. Stanton and J. I. Allston, editors, *Congressus Numerantium*, pages 167–180. Utilitas Mathematica Publishing Inc., Winnepeg, 110 edition, 1995.

3. Robert J. Aumann and Michael Maschler. *Repeated Games with Incomplete Information.* MIT Press, 1995.

4. V. J. Baston and F. A. Bostock. Deception games. *Int. J. Game Theory*, 17(2):129–134, June 1988.

5. V. J. Baston and F. A. Bostock. A simple cover-up game. *Am. Math. Mon.*, 95(9):850, November 1988.

6. Michael Bennett and Russell Richardson Vane, III. Using hypergames for deception planning and counterdeception analysis. *Defense Intelligence Journal*, 15(2):117–138, 2006.

7. Peter Bennett. Toward a theory of hypergames. *Omega*, 5(6):749–751, 1977.

8. Peter Bennett. Bidders and dispenser: manipulative hypergames in a multinational context. *European Journal of Operational Research*, 4(5):293–306, 1980.

9. Peter Bennett and M. Dando. Complex strategic analysis: A hypergame study of the fall of france. *Journal of the Operational Research Society*, 30(1):23–32, 1979.

10. Peter Bennett and Chris S. Huxham. Hypergames and what they do: A "soft O.R." approach. *Journal of the Operational Research Society*, 33(1):41–50, 1982.

11. Peter Bennett, Chris S. Huxham, and M. Dando. Shipping in crisis: A trial run for a "live" application of the hypergame approach. *Omega*, 9(6):579–594, 1981.

12. Emile Borel. The theory of play and integral equations with skew symmetric kernels. *Econometrica*, 21(1):97–100, 1953.

13. Brett J. Borghetti. The environment value of an opponent model. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 40(3):623–633, 2010.

14. Steven J. Brams. Deception in 2x2 games. *J. Peace Sci.*, 2:171–203, 1977.

15. Steven J. Brams and Frank C. Zagare. Double deception: Two against one in three-person games. *Theory and Decision*, 13(1):81–90.

16. Steven J. Brams and Frank C. Zagare. Deception in simple voting games. *Soc. Sci. Res.*, 6:257–272, 1977.

17. R. W. Burns. Deception, technology and the D-day invasion. *Engineering Science and Education Journal*, 4(April):81, 1995.

18. Thomas E. Carroll and Daniel Grosu. A game theoretic investigation of deception in network security. *Security and Communication Networks*, 4:1162–1172, 2011.

19. David A. Castañón, Meir Pachter, and Phillip R. Chandler. A game of deception. *Proc CDC Decis. Control 43rd IEEE Conf.*, 4(4):3364—-3369 Vol.4, 2004.

20. In-Koo Cho and David M. Kreps. Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102(2):179–222, 1987.

21. Austin Lee Davis and Brett J. Borghetti. Deception in game theory and its extensions: a survey. *IEEE Surveys & Tutorials.* to be submitted.

22. Juan J Durillo, José García-Nieto, Antonio J. Nebro, Carlos A. Coello Coello, Francisco Luna, and Enrique Alba. Multi-objective particle swarm optimizers: An experimental comparison. In Matthias Ehrgott, Carlos M Fonseca, Xavier Gandibleux, Jin-Kao Hao, and Marc Sevaux, editors, *Evolutionary Multi-Criterion Optimization. 5th International Conference, EMO 2009*, pages 495–509. Springer. Lecture Notes in Computer Science Vol. 5467, Nantes, France, April 2009.

23. Christopher Elsaesser and Frank J. Stech. Detecting deception by analysis of competing hypotheses. In Alexander Kott and William McEneaney, editors, *Adversarial Reasoning: Computational Approaches to Reading the Opponent's Mind.* Chapman & Hall/CRC, 2007.

24. Andries P. Engelbrecht. *Computational Intelligence: An Introduction.* John Wiley & Sons, Ltd, Hoboken, NJ, 2 edition, 2003.

25. Dario Floreano, Sara Mitri, Stéphane Magnenat, and Laurent Keller. Evolutionary conditions for the emergence of communication in robots. *Curr. Biol.*, 17(6):514–519, March 2007.

26. Niall M. Fraser and Keith W. Hipel. Solving complex conflicts. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(12):805–816, 1979.

27. Niall M. Fraser and Keith W. Hipel. Conflict analysis and bargaining. In *Proceedings; International Conference on Cybernetics and Society*, pages 225–229, 1980.

28. Niall M. Fraser and Keith W. Hipel. Computer assistance in labor-management negotiations. *Interfaces*, 11(2):22–30, April 1981.

29. Niall M. Fraser, Keith W. Hipel, and Jose R. del Monte. Approaches to conflict modeling: A study of a possible USA-USSR nuclear confrontation. *Journal of Policy Modeling*, 5(3):397–417, 1983.

30. Bert Fristedt. The deceptive number changing game, in the absence of symmetry. *Int. J. Game Theory*, 26(2):183–191, 1997.

31. Nandan Garg and Daniel Grosu. Deception in honeynets: A game-theoretic analysis. In *Proceedings of the 2007 IEEE Workshop on Information Assurance, IAW*, pages 107–113, 2007.

32. Bahman Gharesifard and Jorge Cortés. Evolution of players' misperceptions in hypergames under perfect observations. *IEEE Transactions on Automatic Control*, 57(7):1627–1640, 2012.

33. Bahman Gharesifard and Jorge Cortés. Stealthy deception in hypergames under informational asymmetry. *IEEE Trans. Syst. Man, Cybern. Syst.*, 44(6):785–795, 2014.

34. David Hadka and Patrick Reed. Diagnostic assessment of search controls and failure modes in many-objective evolutionary optimization. *Evolutionary Computation*, 20(3):423–452, September 2012.

35. S. N. Hamilton, W. L. Miller, Allen Ott, and O. S. Saydjari. The role of game theory in information warfare. *4th Information survivability . . .*, pages 1–4, 2002.

36. John C. Harsanyi. Games with incomplete information played by 'Bayesian' players, part III. the basic probability distribution of the game. *Manage. Sci.*, 14(7):486–502, 1968.

37. J. P. Hespanha and Y. S. Ateskan. Deception in non-cooperative games with partial information. In *2nd DARPA-JFACC Symposium on Advance in Eterprise Control*, July 2000.

38. J. P. Hespanha and Maria Prandini. Nash equilibria in partial-information games on markov chains. *Proceedings of the 40th IEEE Conference on Decision and Control (Cat. No.01CH37228)*, 3:2102–2107, 2001.

39. João Pedro Hespanha. Application and value of deception. In Alex Kott McEneane and William M., editors, *Adversarial Reasoning: Computational Approaches to Reading the Opponents Mind*, pages 1–26. Taylor and Francis publishing house, 2006.

40. Richards J. Heuer, Jr. Improving intelligence analysis. *Psychol. Intell.*, page 173, 1999.

41. Keith W. Hipel, Muhong Wang, and Niall M. Fraser. Hypergame analysis of the falkland-malvinas conflict. *International Studies Quarterly*, 32:335–358, 1988.

42. Dinh Thai Hoang, Xiao Lu, Dusit Niyato, Ping Wang, Dong In Kim, and Zhu Han. Applications of repeated games in wireless networks: A survey. *Communications Surveys Tutorials, IEEE*, 17(4):2102–2135, Fourthquarter 2015.

43. James Thomas House and George Cybenko. Hypergame theory applied to cyber attack and defense. volume 7666, pages 766604–766604–11, 2010.

44. Eitan Israeli. Sowing doubt optimally in two-person repeated games. *Games and Economic Behavior*, 28(2):203–216, August 1999.

45. Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An introduction to statistical learning*, volume 112. Springer, 2013.

46. J. Kennedy and R. Eberhart. Particle swarm optimization. In *Neural Networks, 1995. Proceedings., IEEE International Conference on*, volume 4, pages 1942–1948 vol.4, November 1995.

47. J. Kennedy and R. Eberhart. *Swarm Intelligence*. Morgan Kaufmann Publishers, San Fransisco, California, 2001.

48. J. Kennedy and R. C. Eberhart. A discrete binary version of the particle swarm algorithm. In *Systems, Man, and Cybernetics, 1997. Computational Cybernetics and Simulation., 1997 IEEE International Conference on*, volume 5, pages 4104–4108 vol.5, October 1997.

49. Keiko Krahnke and Isaac Wanasika. Minimizing strategic deception through individual values. *Journal of Academic and Business Ethics*, 4, 2011.

50. K. T. Lee and K. L. Teo. A game with distorted information. *Naval Research Logistics*, 40:993–1001, 1993.

51. King-Tak Lee. On a deception game with three boxes. *Int. J. Game Theory*, 22(2):89–95, June 1993.

52. Viliam Lisý, Roie Zivan, Katia Sycara, and Michal Pěchouček. Deception in networks of mobile sensing agents. In Michael Luck and Sandip Sen, editors, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, pages 1031–1038, Toronto, 2010.

53. Arthur S. Locke. *Guidance*. Principles of Guided Missile Design *(Merrill, G., Series Editor)*. van Nostrand, New York, 1958.

54. Chris Y. T. Ma, David K. Y. Yau, Xin Lou, and Nageswara S. V. Rao. Markov game analysis for attack-defense of power networks under possible misinformation. *IEEE Transactions on Power Systems*, 28(2):1676–1686, 2013.

55. Akiko Mizutani, Javaan S. Chahl, and Mandyam V. Srinivasan. Insect behaviour: Motion camouflage in dragonflies. *Nature*, 423(6940):604, June 2003.

56. Roger B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, 1991.

57. John Nash. Non-cooperative games. *The Annals of Mathematics, Second Series*, 54(2):286–295, 1951.

58. Antonio J. Nebro, Juan J. Durillo, Jose Garcia-Nieto, Carlos A. Coello Coello, Francisco Luna, and E. Alba. SMPSO: A new pso-based metaheuristic for multi-objective optimization. In *2009 IEEE Symposium on Computational Intelligence in Multi-Criteria Decision-Making (MCDM'2009)*, pages 66–73, Nashville, TN, USA, March 2009. IEEE Press.

59. Animesh Patcha and Jung-Min Park. A Game Theoretic Formulation for Intrusion Detection in Mobile Ad Hoc Networks. *IJ Network Security*, 2(2):131–137, 2006.

60. Martin Peterson. *An Introduction to Decision Theory.* Cambridge University Press, Cambridge, 2009.

61. Radek Píbil, Viliam Lisý, Christopher Kiekintveld, Branislav Bošanský, and Michal Pěchouček. Game theoretic model of strategic honeypot allocation in computer networks. In Jens Grossklags and Jean Walrand, editors, *Decision and Game Theory for Security: Third International Conference, GameSec 2012, Budapest, Hungary, November 5-6, 2012. Proceedings*, pages 201–220. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

62. Ponemon Institute LLC. Intelligence & Incident Response: A Study of U. S. & EMEA Organizations. (February), 2014.

63. J. Poropudas and K. Virtanen. Game-theoretic validation and analysis of air combat simulation models. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 40(5):1057–1070, September 2010.

64. Howard Poston. *Generation of Strategies for Environmental Deception in Two-Player Normal-Form Games.* Thesis, Air Force Institute of Technology, 2015.

65. Marcus Robinson. College of engineering network disabled in response to sophisticated cyberattack. http://news.psu.edu/story/357656/2015/05/15/administration/college-engineering-network-disabled-response-sophisticated, 2015. Accessed: 2015-06-13.

66. G Rosenberg. Enumeration of all extreme equilibria of bimatrix games with integer pivoting and improved degeneracy check. *CDAM Research Report LSE-CDAM-2005-18, London School of Economics*, 42(2010), 2005.

67. A. K. Said and D. A. Hartley. A hypergame approach to crisis decision-making: The 1973 middle east war. *Journal of the Operational Research Society*, 33(10):937–948, 1982.

68. Minoru Sakaguchi. A simple two-player two-sided games of deception. *Sci. Math. Jpn. Online*, 65(1):1045–1053, 2006.

69. Yasuo Sasaki and Kyoichi Kijima. Hypergames and bayesian games: A theoretical comparison of the models of games with incomplete information. *Journal of Systems Science and Complexity*, 25(4):720–735, 2012.

70. Michael C. Shupe, William M. Wright, Keith W. Hipel, and Niall M. Fraser. Nationalization of the suez canal: A hypergame analysis. *Conflict Resolution*, 24(3):477–493, 1980.

71. David Sklansky. *The Theory of Poker.* Two Plus Two Publishing, Las Vegas, NV, fourth edi edition, 2005.

72. David Sklansky and Mason Malmuth. *Hold 'em Poker: For Advanced Players.* Two Plus Two Publishing, Las Vegas, NV, third edit edition, 1999.

73. J. Spencer. A deception game. *Am. Math. Mon.*, 80(4):416–417, April 1973.

74. Frank Stech, Kristin E. Heckman, Phil Hilliard, and Janice R. Ballo. Scientometrics of deception, counter-deception, and deception detection in cyber-space. *PsychNology Journal*, 9(2):79–112, 2011.

75. Masao Allyn Takahashi, Niall M. Fraser, and Keith W. Hipel. A procedure for analyzing hypergames. *European Journal of Operational Research*, 18(1):111–122, October 1984.

76. John W. Thibaut and Harold H. Kelley. *Interpersonal Relations: A Theory of Interdependence.* John Wiley & Sons Inc, New York, New York, USA, 1978.

77. US Dept. of Defense. Joint doctrine for military deception, 2012.

78. Russel R. Vane, III. Hypergame theory for dtgt agents. *American Associate for Artificial Intelligence*, pages 163–165, 2000.

79. Russell Richardson Vane, III. *Using Hypergames to Select Plans in Competitive Environments.* Doctoral dissertation, George Mason University, 2000.

80. Eric Verzuh. *The fast forward MBA in project management.* John Wiley & Sons Inc, Hoboken, NJ, 3 edition, 2008.

81. Carl von Clausewitz. *On War.* Princeton University Press, 1976.

82. John von Neumann. Zur theorie der gesellschaftsspiele. *Math. Ann.*, 100(1):295–320, December 1928.

83. John von Neumann, Oskar Morgenstern, Harold William Kuhn, and Ariel Rubinstein. *Theory of Games and Economic Behavior.* Princon University Press, princeton edition, 2007.

84. Alan R. Wagner and Ronald C. Arkin. Analyzing social situations for human–robot interaction. *Interaction Studies*, 9:277–300, 2008.

85. Alan R. Wagner and Ronald C. Arkin. Robot deception: Recognizing when a robot should deceive. In *Proceedings of IEEE International Symposium on Computational Intelligence in Robotics and Automation, CIRA*, pages 46–54, Daejeon, 2009. IEEE.

86. Alan R. Wagner and Ronald C. Arkin. Acting deceptively: Providing robots with the capacity for deception. *International Journal of Social Robotics*, 3:5–26, 2011.

87. Muhong Wang and K. W. Hipel. Modeling misperceptions in the Persian Gulf crisis. In *Conference Proceedings 1991 IEEE International Conference on Systems, Man, and Cybernetics*, pages 1989–1995, Charlottesville, VA, 1991. IEEE.

88. A. Whiten and R. W. Byrne. Tactical deception in primates. *Behav. Brain Sci.*, 11(02):233, June 1988.

89. D. H. Wolpert and W. G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary COmputation1*, 1(1):67–82, 1997.

90. Yaakov Yavin. Pursuit-evasion differential games with deception or interrupted observation. *Computers & Mathematics with Applications*, 13(1-3):191–203, 1987.

91. Frank C. Zagare. Errata: The Geneva Conference of 1954: A case of tacit deception. *Int. Stud. Q.*, 23(4):600, December 1979.

92. Jun Zhuang, Vicki M. Bier, and Oguzhan Alagoz. Modeling secrecy and deception in a multiple-period attacker-defender signaling game. *European Journal of Operational Research*, 203(2):409–418, 2010.

93. Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evol. Comput.*, 8(2):173–195, June 2000.

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704–0188*

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704–0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202–4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE *(DD–MM–YYYY)* 24–03–2016	**2. REPORT TYPE** Master's Thesis	**3. DATES COVERED** *(From — To)* Jun 2014 — Mar 2016

4. TITLE AND SUBTITLE

Deception in Game Theory:
A Survey and Multiobjective Model

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Davis, Austin, L, Capt, USAF

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Air Force Institute of Technology
Graduate School of Engineering and Management (AFIT/EN)
2950 Hobson Way
WPAFB OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER

AFIT-ENG-MS-16-M-011

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Intentionally left blank

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A:
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

13. SUPPLEMENTARY NOTES

This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

14. ABSTRACT

Game theory is the study of mathematical models of conflict. It provides tools for analyzing dynamic interactions between multiple agents and (in some cases) across multiple interactions. This thesis contains two scholarly articles.

The first article is a survey of game-theoretic models of deception. The survey describes the ways researchers use game theory to measure the practicality of deception, model the mechanisms for performing deception, analyze the outcomes of deception, and respond to, or mitigate the effects of deception. The survey highlights several gaps in the literature. One important gap concerns the benefit-cost-risk trade-off made during deception planning.

To address this research gap, the second article introduces a novel approach for modeling these trade-offs. The approach uses a game theoretic model of deception to define a new multiobjective optimization problem called the deception design problem (DDP). Solutions to the DDP provide courses of deceptive action that are efficient in terms of their benefit, cost, and risk to the deceiver. A case study based on the output of an air-to-air combat simulator demonstrates the DDP in a 7 × 7 normal form game. This approach is the first to evaluate benefit, cost, and risk in a single game theoretic model of deception.

15. SUBJECT TERMS

Game Theory, Deception, Counterdeception, Multiobjective Optimization, Survey

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Dr. B. J. Borghetti, AFIT/ENG
U	U	U	U	90	**19b. TELEPHONE NUMBER** *(include area code)* 937-255-3636x4612 brett.borghetti@afit.edu