APPLIED GAME THEORY? COMPUTATIONAL TECHNIQUES TO OPERATIONALIZE COMPLEX GAMES

by

Michael Macgregor Perry
A Dissertation
Submitted to the
Graduate Faculty
of
George Mason University
in Partial Fulfillment of
The Requirements for the Degree
of
Doctor of Philosophy
Systems Engineering and Operations Research

Committee:

Dr. Hadi El-Amine, Dissertation Director

Dr. Andrew Loerch, Committee Member

Dr. A. Trevor Thrall, Committee Member

Dr. Jie Xu, Committee Member

Dr. John Shortle, Department Chair

Dr. Kenneth S. Ball, Dean, College of Engineering and Computing

Date: 4 NOV 2021

Fall Semester 2021
George Mason University
Fairfax, VA

Applied Game Theory? Computational Techniques to Operationalize Complex Games.

A Dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

by

Michael Macgregor Perry
Master of Arts
Naval War College, 2020
Master of Financial Engineering
University of California, Berkeley, 2015
Bachelor of Arts
University of British Columbia, 2007

Director: Hadi El-Amine, Assistant Professor
Department of Systems Engineering and Operations Research

Fall Semester 2021
George Mason University
Fairfax, VA

DEDICATION

To the memory of Daniel J. Brick.

ACKNOWLEDGEMENTS

I thank my committee, teachers, and classmates. Your insights have been invaluable.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

Approximate probability of correct selection	APCS
Adversarial risk analysis	ARA
Design of experiment	DOE
East China Sea	ECS
Exclusive economic zone	EEZ
Food and Agricultural Organization	FAO
Improvised explosive device	IED
Monitoring, controls, and surveillance	MCS
Optimal computing budget allocation	OCBA
Probability of correct selection	PCS
Response surface methodology	RSM
South China Sea	SCS
Southeast Asian Fisheries Development Center	SEAFDEC
Subgame equilibrium	SGE
Subgame variable	SGV
United Nations	U.N.
United States dollar	USD

ABSTRACT

APPLIED GAME THEORY? COMPUTATIONAL TECHNIQUES TO OPERATIONALIZE COMPLEX GAMES.

Michael Macgregor Perry, Ph.D.

George Mason University, 2021

Dissertation Director: Dr. Hadi El-Amine

Game theory, the mathematical study of strategic interaction, is often criticized as lacking practical application. Such criticism has even come from its most prominent theorists. This dissertation takes the position that these critiques are largely a result of the complexities inherent in game-theoretic analysis, which in turn have relegated most of the literature to the study of over-simplified models, models that are too small in scale for practical application, or both. In light of this, cutting-edge techniques drawn from the operations research literature such as efficient sample allocation, response surface methodologies, robust analysis, and nonconvex binary optimization will be integrated into realistic game theory models. These and other modeling techniques have been slow to integrate into the game theory literature, likely due to the unique challenges games pose as multi-agent optimization problems, and this dissertation thus represents a step forward in operationalizing complex games. Realistic examples will be drawn from

maritime law enforcement and it will be seen that the models presented here have the ability to both explain observed phenomena and contribute to policy development.

1 INTRODUCTION

Game theory attempts to apply mathematical models in situations where at least two intelligent actors (generally, humans) are making decisions. “Intelligence” implies consideration for how others may behave, and game theory thus provides a mathematical framework for analyzing a host situations where competing interests are at play. Of course, a framework in-and-of-itself is not useful. Game theory’s value depends on the ability to state, in mathematical terms, the complex decision-making processes of intelligent actors. If one’s to apply game theory to complex policy issues, which for the purposes of this dissertation can be defined as issues where multiple objectives, nonlinearities, and ambiguity are all ubiquitous, the game is only as useful as it’s ability to capture these things.

Since game theory was introduced as a distinct discipline of study in the 1920s, its success in informing complex policy has been modest. It’s reasonable to assume this is a consequence of the failure to overcome the modeling complexities mentioned above. While scholars have developed the theoretical literature in great detail, deriving analytical solutions in a wide variety of domains such as counterterrorism, contests over natural resources, and civil revolt, these models’ restrictive assumptions have been so far removed from the realities of the complex problems they studied that even the most ardent supporter of quantitative modeling couldn't base a decision on their results. Scholars seem content with this arrangement and analytical solutions remain the dominant form of analysis, perhaps with the less ambitious aim of identifying some directionally correct behavior not otherwise obvious. Significant efforts have, however, been made to use less restrictive assumptions, which has pushed game-theoretic analysis into the realm of computational techniques. While computing power has increased, the complexity and size of the games under study has remained quite small due to an inherent computational challenge of games: the need to solve (at least) two layers of optimizations. In any game, at minimum each intelligent actor must solve an optimization which optimally responds to the actions of all other players; then, an additional optimization mechanism must seek a collection of strategies where each player is optimally responding to all others. The predominance of analytical models with restrictive assumptions, and the modesty of problems addressed via computational techniques, has left game theory on the fringe of policy analysis. While the game-theoretic lexicon has become familiar to most working in business and public policy, where terms like "(non)-zero-sum" have a common meaning and people might routinely ask themselves "what do I think, that he thinks, that I think, …," these merely serve as useful reminders to think about fundamental concepts when dealing with intelligent actors.

This dissertation proposes that game theory can have a more direct contribution to complex policy analysis by incorporating sophisticated computational techniques to analyze models of sufficient complexity to be useful for decision makers. Well-established methods drawn from the operations research literature such as efficient sample allocation, response surface methodologies, and nonconvex binary optimization will be employed to analyze levels of complexity not yet explored in the game theory literature. Recognizing no model, however complex, can ever be accurate enough to produce the one “correct” answer for a complex policy issue, stochastic and robust techniques will be stressed. These techniques, while common in the operations research literature, have been slow to integrate into game theoretic analysis. This is perhaps due to the aforementioned issue of multi-layered optimizations which makes their application less straightforward than in other domains. Taking a practical approach, the dissertation will consider three pertinent examples related to maritime law enforcement to show how advanced methods can both explain observed behavior and inform policy decisions.

The dissertation is structured as follows. The next chapter provides an overarching literature review on the development of game theory and its limited progress analyzing complex policy problems. Three subsequent, stand-alone chapters will analyze particular problems and provide more extensive literature reviews tailored to those problems. The first of these is a generic defend-attack game, as one might encounter when defending oil rigs from robbery. This is a concern, for instance, in the Gulf of Mexico where drug cartels seek to extract fuel from the rigs. It’s conceivably a problem that will plague the South China Sea (SCS) in the future, where territorial disputes have made oil an issue. The primary methodological tool used to analyze the game is statistical selection for discrete choices with stochastic outcomes; in particular, a nested variant of optimal computing budget allocation (OCBA) will be used. The next example deals with the fisheries dispute in the SCS and the need for states to strategically allocate maritime patrols. Tools used include a response surface methodology capable of handling a highly nonlinear model, global optimization to pick sample points based on the response surface, and a robust formulation to account for imperfections in the model. Lastly, an example analyzing the East China Sea (ECS) is analyzed. A key distinction between the SCS and ECS is that in the latter, fishing rights have been agreed to among the major parties, yet illegal encroachments of one nation’s fishermen into the waters of another persist in large numbers. A nonconvex binary program is developed to explain this behavior and make recommendations to resolve it. In the concluding chapter, remarks are given to summarize the work and discuss opportunities for future analysis.

2 A BRIEF LITERATURE REVIEW OF GAMES IN COMPLEX POLICY ANALYSIS

Strategic interaction between actors with competing interests has been a point of analysis for as long as scholars have studyied complex human phenomena; that is to say, since the dawn of scholarship. Mathematical models of such situations can be seen as far back as the eigtheenth centruy (Bellhouse and Fillion 2015). The application of mathematics to strategic interaction didn't explode, however, until John von Neumann published "On the Theory of Games of Strategy" in 1928, followed by the related book 16 years later with Oskar Morgenstern "Theory of Games and Economic Behavior" (von Neumann 1928; von Neumann and Morgenstern 1944). These publications coined the term "game theory" and created a discipline of study that spanned mathematics and the social sciences.

It's useful to consider the development of game theory along two strands: increasing mathematical sophistication on the one hand, and a general infusion of the concepts of strategic interaction into the thinking of qualitative analysis on the other. On the mathematical strand, aside from galvanizing strategic interaction into a sub-discipline, von Neumann's 1928 paper proved the existence of a minimax solution in two-person zero-sum games with perfect information (von Neumann 1928). "Perfect information" implies outcomes are certain and that players know what all others are thinking. John Nash extended this to the case of n-person, non-zero-sum games (Nash 1951). The first breakthrough in relaxing the assumption players know what all others are thinking was Harsanyi's work on what have come to be known as Bayesian games (Harsanyi 1967). A cumbersome limitation of Harsanyi's original work is that players must hold common beliefs about one another. That is, if Player 1 has a particular belief about Player 2 (expressed as a probability distribution for Player 2's objectives), then Player 3 must hold that same belief. This is the so-called theory of "types." It was relaxed by Kadane and Larkey who take a decision-theoretic approach and define probability distributions for adversaries' behavior; no assumption on common beliefs is required (Kadane and Larkey 1982). The adversaraial risk analysis (ARA) framework formalized a method for eliciting those probability distributions (Rios Insua, Rios, and Banks 2009; Banks, Rios, and Rios Insua 2016).

In addition to the above developments, which added realism by modeling the information available to players when they make their decisions, the context in which players make decisions has also been modeled with increasing complexity. Differential games, for example, account for situations where players make several decisions over time and discount future playoffs (Isaacs 1999). Modeling constructs tailored to large populations of actors include evolutionary games (Smith and Price 1973) and mean field games (Jovanovic and Rosenthal 1988). Distinguishing between simultaneous and sequential moves by the players is an important distinction throughout the literature.

As game theory grew in prominance (in academia), it also came under increased scrutiny regarding whether the assumptions it poses about human behavior are legitimate. Games have been criticized for treating intelligent actors both too simplisitcally, not accounting for key decision-making factors such as altruisim and vindictiveness, and for giving humans too much credit in their ability to solve complex optimization problems. Behavioral game theory studies the different ways people make decisions, generally by conducting experiments where subjects play simple games (Camerer 2003). A more mathematical treatment of the issue of uncertainty in player behavior is to take a robust approach that covers all reasonable behaviors. The theory of robust games has been developed in Aghassi and Bertsimas (2006), Kardes (2005), and Crespi, Radi, and Rocca (2020).

These developments have spawned a multitude of analytically solved models applied to various policy issues. A small sampling of applications includes trade wars (Harrison and Rutström 1991), military deterrence (Sorokin 1994), natural resource disputes (Fischer and Mirman 1996; Acemoglu et al. 2012), predicting civil revolt (Chwe 2000; Kiss, Rodríguez-Lara, and Rosa-García 2017), topics in terrorism and insurgency (E. B. De Mesquita 2005; 2007; 2010; S. Wang and Banks 2011; Berman, Shapiro, and Felter 2011), and strategic resource allocation (Gross and Wagner 1950; Kovenock and Roberson 2012a; 2012b; Bier, Oliveros, and Samuelson 2007). All these applications can fairly be described as toy models, and further complexity would make them analytically intractable. While less mainstream, serious developments in applying computational techniques to intractable games have been made. Mere computing power alone has been valuable, such as in generating Monte Carlo samples from stochastic variables describing an adversary's uncertain values and capabilities, in order to elicit an empirical distribution for the adversary's behavior (Rios Insua, Rios, and Banks 2009; Banks, Rios, and Rios Insua 2016). Standard metaheuristics seen in the general optimization literature have been applied to games, such as tabu search (Sureka and Wurman 2005), ant colony optimization (Buer, Homberger, and Gehring 2013), and replicator dynamics (Golman and Page 2009). Algorithms tailored for specific problem classes have also been developed (Nisan et al. 2007). Despite the best computing power, these computational approaches have remained limited to games that are quite small or have easily computable objective functions (more often than not, both conditions are true). Recall games pose a particular computational challenging because, as mentioned in Chapter 1, they always involve multiple layers of optimizations.

The types of problems characterizing complex policy, where highly nonlinear objectives and contraints, many competing objectives, and abiguity requiring stochastic and robust analysis are typically jointly present, have not made much progress in game theory proper. This has relegated game theory's most significnat contributions to the qualitative strand. Terms such as "(non)-zero-sum" have a known meaning and decision makers generally know they must not only consider what an adversary might do, but also what he thinks you may do. This dissertation doesn't seek to provide a historical account of how game theory changed the way qualitative analysts think about complex problems, but a vignette is insightful. In the 1950s, quantiative scholars at RAND (a premier defense think tank) began developing applied models similar to those mentioned above. One scholar was Thomas Schelling, who would go on to author “The Strategy of Conflict” (Schelling 1980). The book introduced the major principles of game theory using plain language and minimal mathematics to analyze several policy issues, largely applicable to Cold War politics. In writing the book, Schelling hoped the theory of games (which is to say the mathematical contributions summarized above) would make its way into the policy world. He’d ultimately have to express his disappointment, noting years later that even the most rudimentary models are still met with resistance. While the book was read and praised by scholars with diverse backgrounds, their intrest was in the general concepts, which Schelling thought were obvious.

Where Schelling was disappointed, other game theorists have acquiesced to the notion that complex phenomena ought not to be modeled mathematically, at least not for the purposes of making specific policy recommendations. The celebrated game theorist Ariel Rubinstein, for example, has stated: “I have not seen, in all my life, a single example where a game theorist could give advice, based on the theory, which was more useful than that of the layman” (Rubinstein, n.d.). As elaborated in his book “Economic Fables,” a game might elucidate a general tendency found in nature, but not a prescriptive policy recommendation (Rubinstein 2012).   Behavioral game theory has also brought into question the value of models, specifically along the lines that the assumed decision-making processes are wrong (Camerer 2003). Player's often exhibit altruism in decision-making, seek to punish bad behavior at their own expense, settle for suboptimal solutions that are good enough (i.e. "satisficing"), exhibit inconsistencies, or simply lack the ability to solve complex optimization problems for their optimal utility. Every one of these objections is an issue of semantics, as these alternative behaviors can be modeled and inserted into a traditional game. The true contribution of behavioral game theory was to force people to think about the variety of decision-making processes people use. The modeling challenge it introduces, assuming the actual decision-making process is unknowable in advance, is how to perform risk analysis around the various possibilities.

This dissertation will respond to the failure of games to make headway into the policy realm by developing solution concepts for games with little-to-no restrictive assumptions. The final determination of whether a model is useful will always remain in the hands of the decision-maker, but it's hoped that emphasizing stochasticity and robustness will provide more confidence in model results. Three specific examples will be given, each with an additional literature review pertinent to the specifics of the problem and methodologies used.

3 COMPUTATIONAL EFFICIENCY IN MULTIVARIATE ADVERSARIAL RISK ANALYSIS MODELS

Note. This chapter was previously published as a stand-alone article in *Decision Analysis* (16 (4), 2019).

This chapter presents a generic defense-attack model with multiple targets. Defense-attack models are not difficult to find in the real world, which explains their popularity in the literature. One example which is consistent with the theme of the later chapters (maritime law enforcement) is the defense of oil rigs from robbery, a common occurrence in the Gulf of Mexico at the hands of drug cartels. It's fair to describe this as a complex problem as oil producers have multiple rigs to protect, real-world objective functions are rarely described by low-order polynomials, and most importantly, the producers lack good information on the cartels' capabilities and values and must therefore perform stochastic analysis. This chapter is distinct from the latter two in that it presents a generic model vice one tailored to a specific situation. This choice was made so that an analytical solution could be derived as a benchmark for the computational technique used. The computational technique, based on optimal computing budget allocation, is seen to converge to the true solution in the majority of cases, produce solutions within 95% of optimality in the remainder, and do so much faster than existing techniques.

3.1. Introduction

Decision making under uncertainty requires the consideration of probable outcomes; at the most rudimentary level a decision maker may assess a prior measure of risk such as $risk = (probability\ event\ occurs) \cdot (disutility\ of\ the\ event)$, and then make a decision with the intent of either reducing the probability of the event, decreasing disutility given the event occurs, or both. In certain disciplines such as engineering this simple approach may be acceptable. Complications arise, however, when the decisions you make are observed by an intelligent adversary whose actions are dependent on your own; in other words, $(probability\ event\ occurs)$ needs to be conditioned on your actions.

While the need to consider the effect of your actions on an adversary's behavior isn't a new idea, in practice the concept is often ignored even at the highest levels of national security (see, e.g. G. G. Brown and Cox Jr. (2011) and Hanley, Jr. (2018)). Game theory provides a natural framework for addressing the problem of decision making in the face of an intelligent adversary and has evolved over the years to address the various shortcomings of its original conception, such as the optimistic assumption that actors know not only how they value various outcomes, but also how their adversaries value outcomes. In brief, von Neumann is often cited as the founder of game theory due to his proof of the existence of a minimax solution in two-person zero sum games with perfect information (von Neumann 1928); subsequent developments were made into solution concepts for more complex problems, the most pertinent in regards to the issue of adversarial uncertainty being the advent of Bayesian games (see Harsanyi (1967)). A more recent development for addressing adversarial uncertainty is a framework that has been termed adversarial risk analysis (ARA). In ARA, decision makers utilize traditional decision-theoretic methodologies while treating their adversary's behavior like any other stochastic parameter, except that the distribution for these parameters is elicited through a prior belief on the distribution of the adversary's utility function and beliefs. General discussions on the merits of the decision-theoretic approach to games can be found in Kadane and Larkey (1982) and Harsanyi, Kadane, and Larkey (1982), as well as in Kadane et al. (2011) and Banks, Rios, and Rios Insua (2016). ARA is a young field and the focus of this chapter is to explore the computational feasibility of ARA models and develop a general-purpose heuristic algorithm for solving them.

In Section 3.3 the details of a general ARA model will be given which will serve as the basis for this chapter. Each a decision maker and an adversary are expected utility maximizers and assumptions will be made about the decision maker's utility function, the adversary's stochastic utility function, and the probability distributions of various events given courses of action of each the decision maker and adversary. Sensible assumptions will be made but the purpose is not to accurately estimate these quantities; the purpose is to assess the computational feasibility of ARA models as they grow in size and complexity.

In the majority of the ARA literature small examples are analyzed where decision makers have limited choices, and utility functions and densities are either binary or easily integrable. In general these conditions will not be met: decision makers typically need to make decisions that are multivariate; outcomes are often more nuanced than a simple binary measure of success or failure; and to accurately model the distribution of outcomes and utilities it would be ideal to have the full spectrum of functions available, regardless of whether they make integration easy. These factors cause exponential growth in the size of the model and to make things more challenging ARA requires the decision maker to assess her adversary's decision-making process, which will also typically be exponentially large. Making ARA models yet more challenging is the fact that the adversary's utility function and beliefs are assumed to be stochastic; this not only means an additional distribution must be integrated over, but there's also typically an elicitation process by which the decision maker assesses her adversary's preferences which involves a computationally expensive Monte Carlo simulation.

In light of these challenges, this chapter addresses the problem of solving an ARA model with a continuous decision space where analytical solutions are unavailable, and a sufficiently fine discretization is used to solve for the optimal strategy. To date, no general-purpose algorithm has been suggested for handling ARA models of this form. Section 3.3 describes such a model, shows how to solve for the optimal decision in exact form (where solutions based on a sufficiently fine discretization are considered exact), and gives descriptive formulae for the computational size of the exact model as a function of the dimensionality of the decision space. The problem blows up quickly, so Section 3.4 describes a simulation-based optimization technique to approximate the model in a feasible amount of time. Section 3.5 presents results of runtimes for variously sized example problems and assesses the statistical confidence in the accuracy the algorithm's output, and also compares it to an alternative methodology employed in the literature, where this chapter's methodology is seen to outperform. Section 3.6 provides an overview of future work that will complement the methodologies described in Section 3.4.

3.2. Literature Review

Rios and Rios Insua (2012) lay out mathematical formulations and give numerical examples for three simple versions of ARA that can be thought of as building blocks for larger models: (i) a simultaneous defend-attack model; (ii) a sequential defend-attack-mitigate model; and (iii) a sequential defend-attack model with private defender information. Various ways their methodology can be expanded to model more detailed scenarios are discussed by Banks, Rios, and Rios Insua (2016) in a comprehensive monograph on ARA. These features include alternative methods to account for level-k thinking (i.e. "he thinks, that I think, that he thinks, that I think …"), multiple adversaries and/or allies, and the expansion of ARA to complex systems that can't be modeled as simple sequences of actions.

In the existing literature on ARA there are a multitude of "toy" examples that have been used to illustrate the above and other concepts while keeping the computational effort low. For example, optimal military convoy routing under the threat of improvised explosive devices (IEDs) was modeled by S. Wang and Banks (2011). Many researchers have studied the optimal allocation of counterterror resources; see, e.g., Rios and Rios Insua (2012), Sevillano, Rios Insua, and Rios (2012), and McLay, Rothschild, and Guikema (2012). Optimal bidding in auctions was studied in Rios Insua, Rios, and Banks (2009). Looking beyond these toy models, the most extensive example found in the literature to date is in Banks, Rios, and Rios Insua (2016), and is based on an actual ARA performed for a client that operates a railway system with multiple stations and is concerned with pickpockets and fare evaders. While the client's decision space is quite large, representing possible security measures, the adversarial decision space is limited (to steal or not to steal), so the problem size remains relatively small. Nevertheless, an exact solution for the optimal security portfolio was computationally impractical so the analysts employed a greedy algorithm with random restarts to search for local optima; in Section 3.5 their approach is applied to the example described in Section 3.3, and the methodology developed in this chapter is shown to outperform.

As noted in Section 3.1, ARA is just one approach to analyzing problems of adversarial conflict. Game theorists in general have studied the above listed applications and as the field evolved analyses accounting for uncertainty in the behavior of other players has become the norm, not the exception. Allocating resources against terrorism and other criminal activities has been studied in C. Wang and Bier (2011), Nikoofal and Zhuang (2012), and Liang and Xiao (2013), to name a few. Protection from IEDs was studied using traditional game theory by Lin and Dayton (2011). Game theory is ripe with other applications that haven't yet been analyzed from an ARA perspective and make for promising future ARA studies; these include models of conventional warfare (e.g. Kovenock and Roberson (2012b)), negotiations among political actors to include fringe figures such as moderate terrorists (e.g. E. B. De Mesquita (2005), Lapan and Sandler (1988), B. B. De Mesquita (1997)), and coordination games among willing participants in revolutionary activities (e.g. Kiss, Rodríguez-Lara, and Rosa-García (2017), Chwe (2000), Edmond (2013)). While the above applications would be interesting to analyze using ARA, meaningful conclusions can only be drawn if methodologies for solving reasonably sized problems quickly are developed.

In their book, Banks, Rios, and Rios Insua (2016) acknowledge computational feasibility will likely be an issue as ARA is developed in greater detail, but that the problem is similar to that faced by other decision theory problems. An example of solving game theory problems with many decision makers and complex sequences of decision can be found in Koller and Milch (2003), where the notion of strategic dominance is used to decompose a problem into several smaller problems that can be solved in sequence. Even in the smaller subproblems the decision space can become exponentially large for the reasons discussed in Section 3.1. If the decision space is a continuous random variable it may be possible to derive analytic solutions to ARA and game theory models, as was done in Zhuang and Bier (2007). However, such results usually rely on closed form expressions for the derivatives of the players' expected utility, which in general won't be possible. This chapter makes the assumption the decision space must be discretized and thus consists of a large yet finite set of possible decisions, each of which is subject to an uncertain outcome due to uncertainty about how an adversary will behave. Statistical selection methods that attempt to select the best among many alternatives have been used extensively in non-adversarial decision problems with uncertainty and generally seek to optimize a measure of confidence the selected alternative is indeed the true optimum. These methods are often grouped by their overarching methodology; popular methodologies include optimal computing budget allocation (OCBA), the expected value of information (e.g. Chick, Branke, and Schmidt (2010)), and sequential elimination methods (e.g., Fan, Hong, and Nelson (2016)). The OCBA methodology first developed by C.-H. Chen et al. (2000) will be used as the overarching framework in this chapter. OCBA has been further developed since its initial conception to account for nuanced, problem-specific applications, and the state-of-the-art methods of OCBA as well as evidence of its sustained superior performance can be found in the textbook by C.-H. Chen and Lee (2011), and a recent paper by Peng et al. (2016).

While powerful, OCBA still requires a preliminary assessment of *all* possible decisions, and thus also may become impractical. Methods combining partitioned based search to identify promising regions of the decision space, and OCBA within the most promising regions, have been explored by, e.g., W. Chen et al. (2014) and Xu et al. (2016). Another method that can be used to identify promising regions of a decision space is that of Bielza, Müller, and Insua (1999), who define an artificial probability distribution for expected utility using the cross section of decision and state variables, which is calibrated via Markov Chain Monte Carlo methods. Determining the marginal distribution in terms of only the decision variables then indicates where in the decision space the optimal decision lies.

3.3. Model Formulation

3.3.1. The Model

This chapter analyzes a two-player sequential game between a decision maker (referred to as "she" throughout this chapter) and her adversary ("he"), where the adversary acts second. As in all ARA models the perspective of the decision maker is taken, and thus her beliefs and utility function are known. In this model, the decision maker implements a strategy, and the adversary then fully observes it and implements his response strategy. Sequential models that continue for multiple moves are obvious extensions of this. Simultaneous models can also be considered an extension, as simultaneous games are often solved using the concept of level-k thinking: a level-1 thinker responds to the expected actions of an actor who doesn't consider his adversary's actions, a level-2 thinker responds to the expected actions of a level-1 thinker, and so on. Players in a simultaneous game are therefore basing their decisions on the expected actions of their adversary, which is mathematically identical to the first mover's problem in a sequential game.

The decision maker and adversary are competing over n distinct targets, which could represent anything from tactical-level concerns such as securing a command center from attack, to high-level strategic concerns such as some measure of economic control over a region. A strategy of the decision maker is denoted as an n-dimensional vector, d, where each element of d represents a particular target towards which she may allocate resources. Similarly, an adversarial strategy a is an n-dimensional vector stating how much resources the adversary invests towards each target. For each target there will be a stochastic level of success achieved with 1 representing the adversary's most favorable outcome and 0 the decision maker's. Denote these levels of success as $S \in \boldsymbol{R^n}$. Given these definitions, the ARA model can now be described.

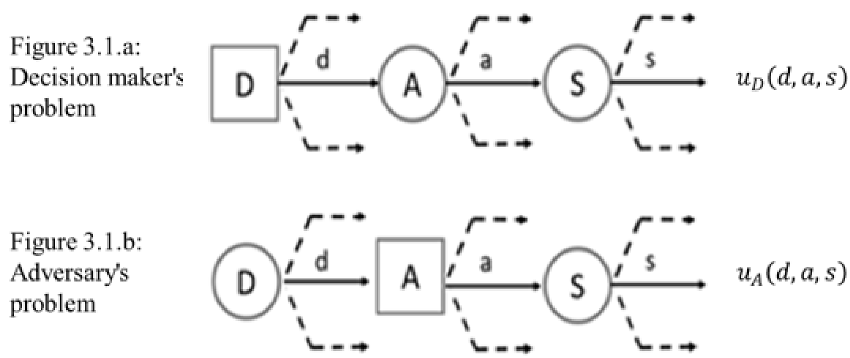

Figure 3.1. Decision trees for the sequential game.

Figures 3.1.a and 3.1.b represents the decisions faced by each the decision maker and adversary, respectively. Circles represent chance nodes where he or she must calculate expected utility given the decisions made up to that point, and squares represent decision nodes. As seen in Figure 3.1.a, the defender can solve her problem of maximizing utility as follows:

D1. For each $(\boldsymbol{d}, \boldsymbol{a})$, find the decision maker's expected utility at chance node $\boldsymbol{S}$:

$\psi_D(d,a) = \int_S u_D(d,a,s) \cdot p_D(s|d,a) ds$, where u_D is the decision maker's utility function and $p_D(s|d,a)$ is her belief about the density of S, given d and a.

D2. For each $\boldsymbol{d}$, find the decision maker's expected utility at chance node $\boldsymbol{A}$:

$\psi_D(d) = \int_A \psi_D(d,a) \cdot p_D(a|d) da$, where A is a random variable representing the adversary's strategy given d and $p_D(a|d)$ is the decision maker's belief about its density.

D3. Maximize expected utility at decision node $\boldsymbol{D}$:

$d^* = \underset{d \in F_D}{\operatorname{argmax}}\, \psi_D(d)$, where F_D is the set of all feasible strategies.

Recall the fundamental principle of ARA that differentiates it from standard game theory: the adversary's utility function and beliefs are uncertain and are elicited using the decision maker's beliefs about her adversary. While the adversary's utility function and density of S are not explicitly present in steps D1 through D3, they materialize implicitly through the decision maker's ability to assess the distribution of a, $p_D(a|d)$, in step D2. Without perfect knowledge of these quantities she's uncertain how her adversary will behave and hence $p_D(a|d)$ is unknown. To overcome this challenge she must make an assumption on the functional form of the attacker's utility function, $U_A(d,a,s,r_u)$, and density of S, $P_A(s|d,a,r_p)$, where r_u and r_p are (possibly multidimensional) realizations of a random variable $R = [R_u \;\; R_p]$ governing the specific forms of U_A and P_A. Also note capital letters have been used to emphasize these are stochastic functions on account of uncertainty in R; this convention will be used throughout the chapter. Making an assumption on the distribution of R, the decision maker can now infer a distribution on her adversary's behavior, $P_D(a|d)$, by analyzing the problem from his perspective. For each possible value of d, the attacker's stochastic behavior can be inferred by tracing Figure 3.1.b backwards from node S to A in the following manner:

A1. Randomly sample a value of R:

Using the assumed distribution of R, the decision maker can randomly sample a single value. Denote this as $r^i = \left[r_u^i \;\; r_p^i\right]$.

A2. For a given $\boldsymbol{d}$, for each $\boldsymbol{a}$ find the adversary's expected utility at chance node $\boldsymbol{S}$:

$\Psi_A^i(d,a) = \int_S U_A\left(d,a,s,r_u^i\right) \cdot P_A\left(s \middle| d,a,r_p^i\right) ds$

A3. Generate a sample of the optimal adversarial strategy at node A, given $\boldsymbol{d}$:

$a^i(d) = \underset{a \in F_A}{\operatorname{argmax}} \Psi_A^i(d,a)$, where F_A is the set of all feasible adversarial strategies.

A4. Infer the density $\boldsymbol{p_D(a|d)}$:

Repeating steps A1 - A3 N_R times generates equally likely samples, $a^1(d), a^2(d), \dots, a^{N_R}(d)$, for the optimal adversarial strategy given d. These samples serve as an estimate the decision maker can use for the density of the adversary's actions, denoted $P_D(a|d)$.

Repeating this process for all d gives empirical densities that can be used in place of $p_D(a|d)$ in step D2, and so the decision maker can now solve her utility maximization problem.

3.3.2. Size of the Model

Analyzing the above model reveals three factors that affect its computational size: the number of feasible strategies, the number of samples of R drawn in step A4, and the precision with which the integrals over S will be estiamted, under the assumption the functional forms of u_D, U_A, $p_D(s)$, $P_A(s)$, and $P_D(a)$ don't offer an analytical solution. The appropriate number of samples to draw from R will depend on its dimensionality, so it will be assumed it contains n elements of uncertainty; recall n is the number of targets being competed over. Denote the number of samples to be drawn in A4 as N_R, and denote the number of intervals to use when discretizing the integrals over each element of S as N_S. Fixing the values $N_R = 10^n$ and $N_S = 10$, the below analysis considers the size of the problem as the number of feasible strategies increases, as this is likely to be the most pertinent parameter to a strategic planner using an ARA model.

The number of feasible strategies will be influenced by two factors: (i) the dimensionality of d and a (i.e. n); and, under the assumption the decision space is in reality continuous but must be discretized for computation, (ii) the discretization used for

d and a. Assume it's sufficient for each element of d and a to take on values in the set $\{0, .1, .2, \dots, .9, 1\}$. The only parameter influencing problem size not yet fixed is n, and analyzing the computational feasibility as n increases will make up the remainder of this chapter.

As a preliminary analysis note that as n increases from 2, to 3, 4, 5, the number of feasible strategies for the decision maker, $|F_D| = \binom{9+n}{n-1}$, increases from 11, to 66, to 286, to 1,001. An identical result obviously holds for $|F_A|$. It can also be shown that the number of integrals over S that must be performed to solve the ARA model is:

$$number\ of\ integrals = |F_D| \cdot |F_A| \cdot (1 + N_R). \tag{3.1}$$

Considering that integrating over S is an n-fold integral to be estimated numerically using N_S intervals per element of S, the actual number of computations required to solve the ARA model is:

$$number\ of\ computations = |F_D| \cdot |F_A| \cdot (1 + N_R) \cdot N_S^n. \tag{3.2}$$

In this chapter four cases for the values of n will be considered: $n \in \{2,3,4,5\}$. The growth in the number of computations in n is enormous, increasing from over 1 million when $n = 2$, to over 4 billion when $n = 3$, 8 trillion when $n = 4$, and finally 10 quadrillion when $n = 5$. When none of the requisite integrals have analytical solutions the model becomes computationally intractable in its exact form when $n = 3$, so as discussed in Sections 3.4 and 3.5 approximation techniques must be used to solve larger problems. Section 3.4 discusses the methods actually employed in this chapter, which take the approach of analyzing the whole decision space in a computationally efficient way. Section 3.6, where future work is discussed, addresses approaches for finding promising regions of the decision space; if identifying these regions can compress the decision space of extremely large problems, then the methods of Section 3.4 can be used to thoroughly assess these relatively small spaces.

3.4. Methodology

The simulation-based methodology that will be used for solving the ARA model of Section 3.3 is based on OCBA, which is a statistical selection technique that's been employed in a variety of contexts to find good solutions quickly. Comparisons will be made for the solve times of the statistical selection method developed here to those of the exact model, and the method's ability to converge to the true optimum will be assessed as well. To facilitate the latter the utility functions, $u_D(d,a,s)$ and $U_A(d,a,s,r_u)$, and the distributions $p_D(s|d,a)$ and $P_A(s|d,a,r_p)$ have been selected so the integrals in steps D1 and A2 give analytical solutions for $\psi_D(d,a)$ and $\Psi_A(d,a)$ and the optimization over the adversary's behavior in A3 can be solved efficiently using standard algorithms, and thus the model can be solved in exact form in a reasonable amount of time. Because in practice the interest is in solving problems without analytical solutions for $\psi_D(d,a)$ and $\Psi_A(d,a)$, and where the optimization for the adversary's behavior is hard, the model will also be solved in exact form by discretizing these integrals and then solving the optimization model in A3 via full enumeration for the purpose of assessing the computational advantages of statistical selection.

Before discussing the details of how OCBA is employed, a brief narrative example for the ARA model to be solved is given as well as the functional forms of its utility and density functions.

3.4.1. A Modified Colonel Blotto Game

Colonel Blotto games are a class of widely studied game theory models where one military commander, Colonel Blotto, must allocate her forces to n distinct battlefields, knowing that her adversary, Colonel Klink, is simultaneously making the same decision. The objective is to win control of the most battlefields where each field is won based on some (possibly stochastic) function of the amount of forces deployed. This chapter modifies the game to assume there's an ongoing military conflict in which Colonel Blotto is considering intervening. If she intervenes to aid her allies, she knows with certainty Colonel Klink will respond by intervening in aid of the other side, but she's uncertain as to which battlefields he values most and hence how he'll respond. Blotto's only uncertainty about Klink is in his utility function, and Blotto and Klink share a common density over the outcome on a battlefield, $p(s|d,a) = p_D(s|d,a) = p_A(s|d,a)$. The modified game is a sequential model: Colonel Klink only acts after observing Colonel Blotto's decision.

In this context, Colonel Blotto is the decision maker and Colonel Klink the adversary. The decision vectors d and a now represent what percentage of forces each Blotto and Klink deploy to each of the n zones, and S is a random vector representing the degree of success in each battlefield ($S_i \in [0,1]$, where $S_i = 1$ is the best possible outcome for Colonel Klink).

3.4.2. Utility and Density Functions

Assume the levels of success on each battlefield follow independent uniform distributions over intervals of length 0.1, $S_i \sim Uni(h_i, h_i + .1)$, where $h_i \in [0, .9]$ depends on the i^{th} elements of d and a. In particular, $h_i(d_i, a_i) = -\frac{2}{4.6}\left(\log\left(c_{A,i}a_i + 1\right) - \log\left(c_{D,i}d_i + 1\right)\right) + C_{H,i}$, where $c_{A,i}, c_{D,i}$, and $C_{H,i}$ are measures of how difficult it is to attack target i. Blotto's and Klink's utility functions are:

$$u_D(s) = \frac{\sum_{i=1}^{n} v_i\left(e^{-4.6\left(s_i - C_{H,i} - .05\right)} - 1\right)}{n}, \tag{3.3}$$

$$U_A(s, r) = \frac{\sum_{i=1}^{n} r_i\left(1 - e^{-4.6\left(s_i - C_{H,i} - .05\right)}\right)}{n}, \tag{3.4}$$

where v_i is the known value Blotto places on battlefield i and r_i is the uncertain value Klink places on battlefield i.

The proof these specifications result in analytic solutions for $\psi_D(d, a)$ and $\Psi_A(d, a)$ is in Appendix 3.A, and the formulation for $a^i(d) = \underset{a \in F_A}{\operatorname{argmax}} \Psi_A^i(d, a)$ as a simple quadratic integer program is in Appendix 3.B. The intuitive appeal of these specifications is that h_i is increasing in a_i and decreasing in d_i, and determines a relatively tight bound of length 0.1 on the possible outcomes of each S_i based on how much effort each Blotto and Klink apply to field i. The rate of change of S_i with respect to a_i and d_i is driven by the (fixed) parameters $c_{A,i}$ and $c_{D,i}$. If each Blotto and Klink invest zero resources on zone i, then $h_i = C_{H,i}$. In the example analyzed in Section 3.5, $C_{H,i}$ will be non-zero, indicative of the fact they're intervening in an ongoing conflict where, without intervention, the S_i values will be non-zero. The utility functions are exponential and if all $S_i = C_{H,i} + .05$, representing no change to the status quo, then utilities evaluate to zero. Values of S_i above $C_{H,i} + .05$ lead to positive utility for Klink proportional to the importance he attaches to battlefield i, r_i, while values below $C_{H,i} + .05$ lead to positive utilities for Blotto proportional to v_i. The "4.6" constant was chosen so that in the extreme case when $C_{H,i} = 0$ and $s_i = 1$, Klink receives utility of approximately $\frac{r_i}{n}$ and Blotto is decremented by $\frac{v_i}{n}$ (because $e^{-4.6} \approx 0$). For the distribution of the unknown parameters R_i of which r_i are realizations, it's assumed each R_i follows a triangular distribution: $R_i \sim triangular(l_i, m_i, u_i)$. In practice, these model parameters would be fit rigorously by codifying intelligence analysts' knowledge into measurable functions and distributions.

3.4.3. Optimal Computing Budget Allocation for Solving Large ARA Models

To solve for Colonel Blotto's optimal decision an algorithm based on the optimal computing budget allocation (OCBA) scheme of C.-H. Chen et al. (2000) is used. OCBA is a general technique to support decision making under uncertainty. While it assumes normality of u_D, theoretical methods that generalize OCBA to any distribution are a current point of research and a few non-normal extensions have been derived in Glynn and Juneja (2004), and C.-H. Chen et al. (2000) showed that in practice OCBA performs well for non-normal utility functions. Empirical results for the example presented in this chapter show highly normal behavior (see Figure 3.3 of Section 3.5.2).

Developed for the purpose of choosing a good strategy under uncertainty when fully evaluating all possibilities is impractical, OCBA evaluates potential strategies through a series of iterations. In the first iteration all strategies are given preliminary examination by taking a small number of samples of the random variables causing the uncertainty (in this case, S and R). In subsequent iterations, OCBA solves a nonlinear optimization problem to maximize an approximation of the probability of selecting the best strategy, subject to a constraint on the total number of samples to be drawn per iteration. The iterative process continues until either sufficient confidence an optimal strategy has been found is achieved, or the total computational budget is exhausted. Formally, during each iteration OCBA's sampling allocation scheme yields the result in Theorem 3.1:

Theorem 3.1 (proved in C.-H. Chen et al. (2000))

Given a finite number of samples to be allocated among k competing strategies, whose utility each follow a normal distribution, $u_i \sim N(\mu_i, \sigma_i^2)$ for $i = 1,2,\ldots,k$, the approximate probability of funding the true optimal strategy (APCS) is asymptotically maximized using the following allocation scheme:

$$\frac{N_i}{N_j} = \left(\frac{\sigma_i/(\mu_b - \mu_i)}{\sigma_j/(\mu_b - \mu_j)}\right)^2, i \neq j \neq b, \tag{3.5}$$

$$N_b = \sigma_b \sqrt{\sum_{i \neq b} \frac{N_i^2}{\sigma_i^2}},$$

$\sum_{i=1}^{k} N_i =$ total sampling budget, and

$b = \underset{i}{\operatorname{argmax}}\, \mu_i$ (i.e. the index for the best strategy, based on the highest mean utility).

When applying Theorem 3.1 in practice, the normal parameters μ and σ are estimated using the most current sample data of the OCBA-based algorithm described above. The approximation for the probability of correctly selecting the true optimum, APCS, is the objective function being maximized and is defined using the Bonferroni (1936) lower bound on the actual probability of correct selection:

$$P(u_b > u_i, \forall\, i \neq b) \geq 1 - \sum_{i \neq b} P(u_i > u_b) = 1 - \sum_{i \neq b} \Phi\left(\frac{\mu_i - \mu_b}{\sqrt{\frac{\sigma_b^2}{N_b} + \frac{\sigma_i^2}{N_i}}}\right) := APCS, \tag{3.6}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function.

The choice of APCS as the metric to be maximized is partly pragmatic, as it allows the optimal sampling problem in Theorem 3.1 to be solved analytically. However, APCS is also an appealing metric from a decision-theoretic standpoint as it's a lower bound on the true probability of correct selection (PCS); lacking a reliable metric for PCS, a decision maker will likely be most interested in knowing how low PCS could be (i.e. APCS). Also, as shown in C.-H. Chen and Lee (2011), APCS converges to PCS under mild conditions and numerical tests show strong performance in identifying strategies with the true PCS. Aside from fixing a total computational budget across iterations, there's no general-purpose stopping criterion that's used in the OCBA literature. A sufficiently high APCS may seem a natural stopping criterion but has issues when two strategies have expected utilities that are close. In ex post analysis, this chapter outlines a simple method for adjusting APCS when two or more expected utilities are nearly indistinguishable, but for a stopping criterion a metric is defined based on a lack of material changes in the most promising strategies from one iteration to the next; this is detailed in Appendix 3.C.

Applying this to the ARA model, samples of Colonel Blotto's objective function must be drawn for all values of $d \in F_D$. This is done by simultaneously drawing samples of S and A. Sampling a value $S = s$ is straightforward as each element is $S_i \sim Uni(h_i, h_i + .1)$, with h_i deterministically in d_i and a_i. The difficulty is in sampling $A = a$ as $p_D(a|d)$ is unknown. To overcome this, a sample is instead drawn for $R = r$ and then transformed into a sample of Colonel Klink's optimal behavior by solving $a^*(d) = argmax_a \int_S U_A(d, a, s, r) \cdot p(s|d, a) ds$. Because this is an n-dimensional integral that may not have an analytical solution, it becomes computationally expensive as n increases. Thus, a nested OCBA is used to evaluate it, where for each $a \in F_A$ samples from S are drawn according to the OCBA sampling allocation rules. Together with the sample s in the higher-level OCBA representing Blotto's problem, this gives a single sample of $u_D(d, a, s)$. Repeating the process per the allocation scheme in (3.5) gives a sample set that can be used to estimate $\psi_D(d)$; repeating for all d allows the allocation scheme (3.5) to be used for the next iteration.

Algorithm 3.1 formalizes the solution technique to be used in Section 3.5. To provide a descriptive summary of the algorithm, for each strategy d, multiple draws are taken from (S, R) and each of these draws requires performing a nested OCBA to transform the sampled value of R into a sampled value of A. The drawn samples are used to estimate the utility of each strategy d, and the process repeats until the stopping criteria in Appendix 3.C is met. This process is referred to as one "trial," as in a Bernoulli trial where a particular strategy either is or is not proposed as the optimal strategy. Trials are ran until a 95% confidence level is reached for the most frequently observed outcome of a single trial being, in fact, the most likely outcome. This confidence level is calculated using Wilson's (1927) Bernoulli confidence interval:

Lemma 3.1 (see proof in Appendix 3.D)

Assume an experiment is being conducted with k possible outcomes, E_1, E_2, …, E_k, so that the occurrence of an event E_i is a Bernoulli random variable, B_i, with success probability p_i, where $\sum_{i=1}^{k} p_i = 1$. Assume all p_i values are unknown, that N trials of the experiment are conducted with $B_{ij} \coloneqq 1$ if the j^{th} trial results in event E_i, and $B_{ij} \coloneqq 0$ otherwise, and define $f_i \coloneqq N - \sum_{j=1}^{N} B_{ij}$, the number of times event E_i did not occur.

A sufficient condition to be $1 - \alpha$ confident the most frequently observed outcome, E_b, is in fact the most likely, is that:

$$.5 < \frac{N - f_b + \frac{z^2}{2}}{N + z^2} - \frac{z}{N + z^2}\sqrt{\frac{f_b(N - f_b)}{N^3} + \frac{z^2}{4}}, \quad (3.7)$$

where $z \coloneqq \Phi^{-1}(1 - \alpha)$.

Formally, the algorithm proceeds as follows:

Algorithm 3.1. Nested OCBA to Optimize the Decision Maker's Strategy.

1. Create an $|F_D|$-dimensional list, B, indicating the number of times each of the decision maker's strategies has been proposed by steps 2 through 18 as an optimum. Initiate all elements to 0.

    2. Create an $|F_D|$-dimensional list, $OCBA_D$, to store samples of $u_D(d, a, s)$ for all d.

    3. Initiate the number of samples for each d in the initial iteration of OCBA to $N_d = N_{init}$, and set the iteration number to $iter_D = 0$.

    4. Increment $iter_D = iter_D + 1$.

    5. For all d:

        6. For all i = 1:N_d:

            7. Generate a sample, r, from R.

            Determine the optimal adversarial strategy for $\boldsymbol{d, r}$ using nested OCBA

            8. Create an $|F_A|$-dimensional list, $OCBA_A$, to store samples of $U_A(d, a, s, r)$ for all a.

9. Initiate the number of samples for each a in the initial iteration of OCBA to $N_a = N_{init}$, and set the iteration number to $iter_A = 0$.

10. Increment $iter_A = iter_A + 1$.

11. For all a:

    12. For all j = 1:N_a:

        13. Generate a sample, s_A, from $S \sim Uni(h(d, a),\ h(d, a) + .1)$.

        14. Calculate $U_A(d, a, s_A, r)$ and store it in $OCBA_A(a)$.

15. If $iter_A = 20$ or the stopping criterion of Appendix 3.C is reached, set the adversary's optimal behavior for this value of (d, r) to $a^*(d, r) \coloneqq$ the adversarial strategy with the highest sample mean in $OCBA_A$. Otherwise, use the OCBA allocation scheme, (3.5), to update N_a for all a, and return to step 10. Note that when calculating (3.5), in the event the sample standard deviation is 0 for some strategy i we set $s_i = .0001$. Values of N are rounded up to get an integer number of samples, effectively setting $N_i = 1$.

16. Generate a sample, s_D, from $S \sim Uni(h\big(d, a^*(d, r)\big),$ $h\big(d, a^*(d, r)\big) + .1)$.

17. Calculate $u_D(d, a^*(d, r), s_D)$ and store it in $OCBA_D(d)$.

18. If $iter_D = 20$ or the stopping criterion of Appendix 5.C is reached, select the strategy d with the highest sample mean in $OCBA_D$, and increment the element of B corresponding to the selected optimum. Otherwise, use the OCBA allocation scheme, (3.5), to update N_d for all d, and return to step 4 (see note in step 15 for the case when $s_i = 0$ for some strategy i).

19. If the most frequently observed strategy in B is in fact the most likely outcome of a trial with at least 95% confidence (see Lemma 3.1), stop; select the most frequently observed strategy in B as the (approximated) optimal solution. Otherwise, repeat from step 2.

Theorem 3.2 provides a guide for estimating the expected number of trials for Algorithm 3.1 to terminate, based on the required statistical confidence in its output:

Theorem 3.2 (see proof in Appendix 3.D)

Assume multiple trials of the experiment described in Lemma 3.1 are to be performed and again define by f_b the number of trials in which the most frequent outcome, E_b, does not occur. Given f_b, denote by $N(f_b)$ the minimum number of trials to be $1 - \alpha$ confident E_b is in fact the most likely outcome of a single trial. Let $n_0 \coloneqq \lceil z^2 \rceil$, where $z \coloneqq \Phi^{-1}(1 - \alpha)$. The following are true:

1. Either $N(f_b) = 2f_b + n_0$, or $N(f_b) = 2f_b + n_0 + 1$.

2. A sufficient condition for $N(f_b) = 2f_b + n_0$ is: $\frac{\sqrt{3}}{18n_0} < \left(\frac{n_0}{2z}\right)^2 - \frac{z^2}{2}$.

3. Assuming the condition in part 2 is met and given p_b, the true probability of the most frequently observed outcome, the expected number of trials until Algorithm 3.1 terminates can be calculated as:

$$E[N(f_b)|p_b] = p_b^2 \cdot \sum_{f_b=0}^{\infty}(2f_b + n_0) \cdot P(!\,N(f_b - 1), f_b|p_b), \qquad (3.8)$$

where $P(!\,N(f_b - j), i|p_b)$, the probability of not terminating after $N(f_b - j)$ trials while observing exactly i failures during those trials, can be calculated recursively as:

$$P(!\,N(f_b - j), i|p_b) = \qquad (3.9)$$

$$\sum_{k=k_{min}}^{k_{max}} P(!\,N(f_b - j - 1), k|p_b) \cdot p_b^{2-i+k} \cdot (1 - p_b)^{i-k} \cdot \binom{2}{i-k},$$

where $k_{min} \coloneqq \max\{f_b - j, i - 2\}$, $k_{max} = \min\{2(f_b - j - 1) + n_0, i\}$, and the recursion can be started using $P(!\,N(0), k|p_b) = \binom{n_0}{k} \cdot p_b^{n_0-k} \cdot (1 - p_b)^k$ for $1 \le k \le n_0$, and 0 otherwise.

Although this is an infinite sum, it converges to its true value quickly as long as p_b is sufficiently higher than 0.5.

The sufficient condition in part 2 of Theorem 3.2 is met for the vast majority of confidence levels, including when using the required level of 95%. While the expected number of trials can't be calculated without knowledge of p_b, a reasonable upper bound can be obtained by assuming any well-constructed algorithm will have $p_b \geq 60\%$; along with 95% required confidence, this gives $E[N(f_b)|.6] = 15$. In Appendix 3.D many trials of the algorithm are ran (more than required to terminate) and its observed empirically that $p_b \approx 76.47\%$; this gives an expected number of trials of 5.7. As a technical note, to deal with situations where the estimated expected utilities of two or more strategies are essentially equivalent, in order to make the events E_i distinct its assumed the most frequently observed among them (in prior trials) is the optimum for the current trial. To avoid path dependencies the most frequently observed is recalculated after each trial. Ties are broken arbitrarily. Such a condition is required since without it two strategies that for all intents and purposes yield equivalent utilities will cause an excessive number of trials to be run in expectation, even once it's become clear the strategies are indistinguishable and either can be selected.

3.5. Discussion of Findings

Four values of n were analyzed, $n = 2, 3, 4$, and 5, with parameters as detailed in the next subsection. For each case Algorithm 3.1 was ran to estimate the optimal solution to Colonel Blotto's problem, the time required to reach that solution was measured, and ex post analysis was performed using the sample utilities generated to calculate a lower bound on the probability the proposed optimum is in fact the true optimum. The exact optimal solution was also computed by exploiting the analytical forms of $\psi_D(d, a)$ and $\Psi_A(d, a)$ and the quadratic structure of $a^i(d) = \underset{a \in F_A}{\operatorname{argmax}} \Psi_A^i(d, a)$, and compared to the solution found by Algorithm 3.1. For the remainder of the chapter this solution method will simply be referred to as the "partial analytic" method. Even using the partial analytic method the computational advantages of Algorithm 3.1 become apparent when $n \geq 4$.

To make clear the computational savings when no analytical solutions are available (as will be the case in general) the time required to solve the exact model when $\psi_D(d,a)$ and $\Psi_A(d,a)$ must be evaluated numerically and $a^i(d) = \underset{a \in F_A}{\operatorname{argmax}} \Psi_A^i(d,a)$ solved via full enumeration was also assessed; this solution method will be referred to as the "fully numeric" method.

3.5.1. Parameter Values and Sampling Budgets

As n is increased from 2 to 5, the below parameter values are used. When $n = 2$, only the first two elements of each parameter are used, when $n = 3$ the first three elements are used, and so on.

$C_H = [.4,\ .35,\ .4,\ .4,\ .3]$,

$c_A = [-.4984, -.4984, -.5529, -.6015, -.6834]$,

$c_D = [-.4984, -.4373, -.4373, -.5529, -.4626]$,

$v = [1.3,\ .8,\ 1.25,\ .7,\ 1.1]$, and

$R = [R_1, R_2, R_3, R_4, R_5]$,
where $R_1 \sim triangular(.8,\ 1,\ 1.5)$, $R_2 \sim triangular(.5,\ .8,\ 2.5)$,

$R_3 \sim triangular(1,\ 1.5,\ 3.5)$, $R_4 \sim triangular(.3,\ .7,\ 1.1)$, and

$R_5 \sim triangular(.6,\ 1.1,\ 1.9)$.

$N_{init} = 2^n$ was used as the number of samples to allocate to each of Colonel Blotto's strategies in the first iteration of Algorithm 3.1. In later iterations, a total of 5^n samples were allocated across the strategies per the OCBA allocation scheme. The same sampling rules were used for the nested OCBAs to solve Colonel Klink's decision problem.

3.5.2. Run Times and Ex Post Analysis

Solve times are now compared for the fully numeric method, the partial analytic method, and Algorithm 3.1. Illustrating the difficulty in solving large ARA problems by brute force, the fully numeric method took 26 minutes when $n = 2$ and was computationally infeasible for $n \geq 3$. The partial analytic method took 1 minute when $n = 2$, 9 minutes when $n = 3$, 597 minutes (10 hours) when $n = 4$, and 62,000 minutes (43 days) when $n = 5$. In contrast, Algorithm 3.1 took 1 minute when $n = 2$ and terminated after 3 trials, 24 minutes when $n = 3$ and terminated after 7 trials, 279 minutes (4.65 hours) when $n = 4$ and terminated after 5 trials, and 2,618 minutes (44 hours) when $n = 5$ and terminated after 3 trials. These results are summarized in Table 3.1. The number or trials required for Algorithm 3.1 to terminate is subject to random variation in the samples that are drawn, and in Appendix 3.E the analysis is repeated many times for $n = 4$ to assess the expected number of trials to termination; in all cases Algorithm 3.1 selects the correct optimum and the expected number of trials is rather low (5.7), but can be as high as 13.

Figure 3.2 shows graphically that Algorithm 3.1 identified the true optimum for all values of n. The true utilities of each strategy are plotted along with those estimated after the first iteration of the first trial and the estimated utilities upon termination. Plotting the entire decision space for $n \geq 3$ would make the plots difficult to read so only the area surrounding the true optimum is shown; no estimated utilities of Algorithm 3.1 (at termination) outside the range plotted exceeded the highest within the plotted range. The important thing to note about Figure 3.2 is that while divergence between the true expected utilities and those computed by Algorithm 3.1 remains in some parts of the decision space, Algorithm 3.1 has converged near the peak values of expected utility, reflecting the computational budgetary allocation scheme that dedicates the majority of samples to understanding the most promising regions of the decision space. As can be seen in Figures 3.2.b and 3.2.c, a false optimum would have been identified had the algorithm stopped after just 1 iteration of the first trial, but given additional iterations the algorithm arrived at the true optimum.

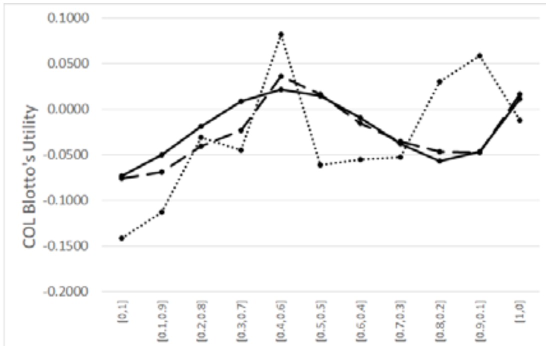

Figure 3.2.a: n = 2

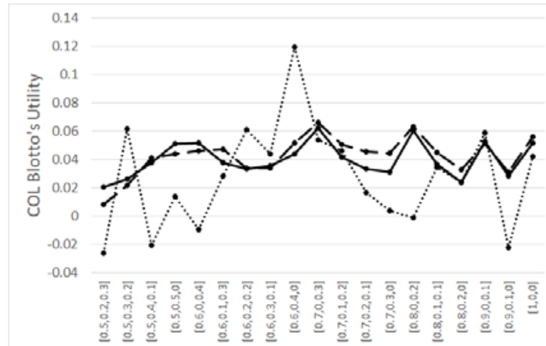

Figure 3.2.b: n = 3, near $d^* = [.7, 0, .3]$

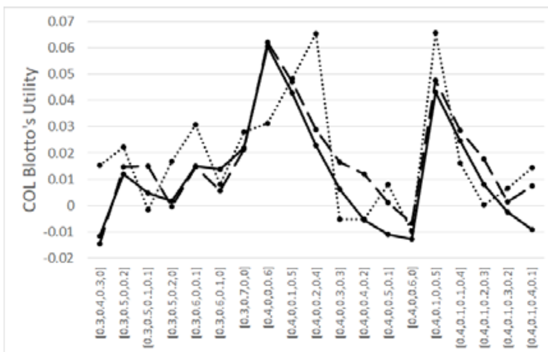

Figure 3.2.c: n = 4, near $d^* = [.4, 0, 0, .6]$

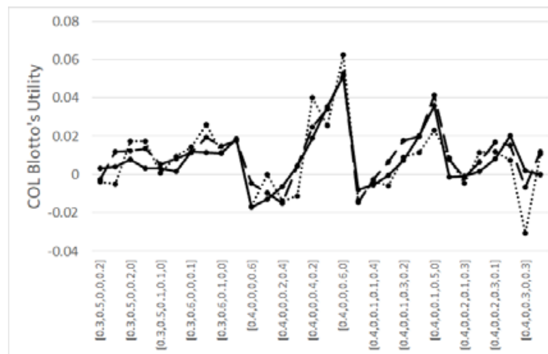

Figure 3.2.d: n = 5, near $d^* = [.4, 0,0, .6, 0]$

Exact Simulation-based Solution, first iteration Simulation-based Solution, terminal

Figure 3.2. Exact values of $\psi_D(d)$ and estimates from Algorithm 3.1

In general an exact solution wouldn't be available to compare the results of Algorithm 3.1 to, so to add confidence in its solution the sample values of utility, $u_D(d)$ for all $d \in F_D$, were used to calculate a lower bound on the probability the strategy with the highest sample mean utility, $d_b|\{\hat{u}_D(d_b) > \hat{u}_D(d_i)\ \forall\ i \neq b\}$, does in fact yield the highest utility. Using the assumption that $u_D(d)$ is normally distributed this is a straightforward calculation using the Bonferonni inequality described in equation (3.6). To justify the normality assumption, normal qq-plots for $u_D(d)$ were created for a variety of values of d and n. Figure 3.3 shows the plots for three particular values of d when $n = 4$ which are illustrative of the results seen for other values; all showed clear normality despite the occasional bit of deviation in the tails (Figure 3.3.c was the most extreme case found). Calculating the lower bounds on the probability the algorithm's proposed optimum is in fact the true optimum (APCS), it was found that: in the case of $n = 2$, APCS evaluated to 91.41%; when $n = 3$, it was 57.01%; when $n = 4$ it's 94.59%; and when $n = 5$ it's 98.39%. These lower bounds are summarized in Table 3.1. The markedly lower value of APCS when $n = 3$ is due to a point noted in Section 3.4: as can be seen in Figure 3.2.b, two strategies, $d_b = [.7, 0, .3]$ and $d_2 = [.8, 0, .2]$, have essentially identical expected utilities of .0660 and .0630, respectively. In the summation to calculate APCS, along with their respective standard errors (.0036 and .0049) the term $P(u_2 > u_b) = 31.12\%$ is subtracted out. If this had not been subtracted, essentially treating strategies d_b and d_2 as one and the same, the lower bound on the probability of correct selection would have been 88.19%. In Appendix 3.E where the analysis is repeated for $n = 4$ several times, a simple calculation to adjust APCS values is used that doesn't require decision maker's to visually inspect expected utilities for virtual equivalency.

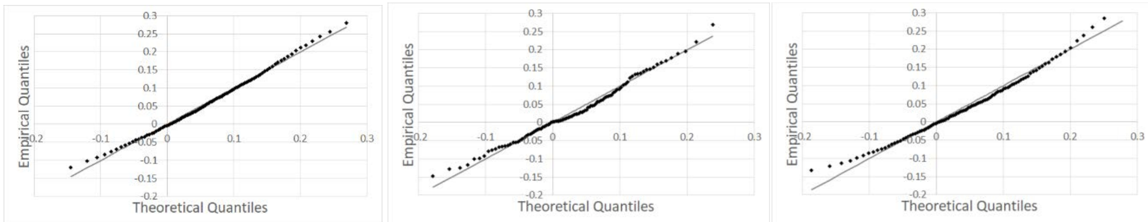

Figure 3.3.a: $d = [.4,0,0,.6]$ *Figure 3.3.b:* $d = [0,0,.2,.8]$ *Figure 3.3.c:* $d = [.9,0,.1,0]$

Figure 3.3. Normal qq-plot for $u_D(d)$ using 1st to 99th quantiles.

Table 3.1. Run times and lower bounds on probability of correct selection.

	n = 2	n = 3	n = 4	n = 5
Exact model solve time, without using analytical integrals over S and QIP	26			
Exact model solve time, using analytical integrals over S and QIP	1	9	597	62,000
Algorithm 1 solve time	1	24	279	2,618
Number of trials	3	7	5	3
Algorithm 1 PCS lower bound	91.41%	57.01% *	94.59%	98.39%

All times measured in minutes

** When the two best strategies are considered equivalent, the bound on PCS is 88.19%*

3.5.3. Benchmarking Against a Greedy Search with Random Restarts

As noted in Sections 3.1 and 3.2, there are few examples in the literature of solution techniques for large ARAs, so the results presented in this chapter are useful in their own right. However, this subsection compares Algorithm 3.1 to the technique used by Banks, Rios, and Rios Insua (2016) to find a heuristic solution to a fairly large ARA model. They used a greedy algorithm with random restarts that begins with a randomly selected feasible strategy, generates a sufficient number of samples to estimate the expected utility of that strategy, and then permutes a single element of the strategy to reach a new feasible strategy. Samples are again drawn, and if this yields a new incumbent solution the algorithm continues with an additional permutation to this new incumbent; otherwise, the algorithm returns to the incumbent and tries an alternative permutation. Applying this to the modified Colonel Blotto example, a nested greedy algorithm is required where for each strategy of Blotto he must draw random samples for Klink's r value, and for each of these samples Blotto must use a greedy algorithm to analyze the decision problem from Klink's perspective. Using the case when $n = 4$, it was assumed 100 samples was sufficient for each iteration of the algorithm; as observed empirically, a single trial of Algorithm 3.1 generally allocates several hundred samples to each of the best few strategies, and 100 samples was seen to be a natural segregator between decent strategies and strategies that were entirely ignored following the initial iteration. As a stopping criterion it’s required that the probability of finding a new local maximum with another random restart be less than 5% as calculated using Laplace's law of succession, which is consistent with Algorithm 3.1's 95% required confidence level. This algorithm took 3,371 minutes (56 hours) to terminate with $n = 4$; as detailed in Appendix 3.E, in expectation Algorithm 3.1 will take just 4.92 hours to terminate when $n = 4$. Further, though the greedy algorithm did find the true optimum, due to the relatively small sample size the APCS was -82.00% (i.e. no better than the trivial lower bound of 0%); while this is only a lower bound on the probability of correctly selecting the optimal solution, it clearly indicates much less confidence than when Algorithm 3.1 is used. Indeed, as shown in Appendix 3.E, even after accounting for virtual equivalency in mean utilities APCS for the greedy algorithm is less than 0%. As an additional benchmark, experiments were ran to find the number of samples to generate for each strategy in the greedy algorithm such that it terminates in about 4.95 hours. The number of samples needed to be reduced to 7 per strategy. Of course, as APCS was meaningless when 100 samples per strategy were drawn, it's likewise meaningless with just 7 samples per strategy (-3,034.75%).

3.5.4. Summary of Results

Algorithm 3.1 accurately estimated $\psi_D(d)$ for all values of n tested, and in Appendix 3.E the analysis is repeated several times for $n = 4$ to assess the accuracy rate in selecting the true optimum. In brief, for the parameters in this section the algorithm always identifies the true optimum in 102 runs of the algorithm, and using three alternative parameter sets and an additional 144 runs it always selects a strategy whose utility is at least 95% of the true optimum, though is typically much closer to 100% of optimality. Equally important to the accuracy of the algorithm is its speed, and the run times summarized in Table 3.1 as well as Appendix 3.E clearly show it makes generically specified ARA problems that would otherwise be intractable solvable in a reasonable amount of time.

3.6. Conclusion and Future Work

This chapter has presented a simulation-based algorithm for a two-player, sequential ARA model that terminates in reasonable time when the decision spaces for each the decision maker and adversary are large and the functional forms of utilities and state variable distributions don’t offer analytical solutions. By running multiple independent trials, as defined by steps 2 through 18 of Algorithm 3.1, the algorithm terminates only after sufficient statistical confidence is achieved. Further, in the example of Section 3.5 ex post analysis yielded lower bounds on the probability the algorithm produced the true optimum that were quite high. The example used to assess the algorithm was a modified Colonel Blotto game and the number of battlefields was used as a variable parameter to increase the size of the model, but the algorithm is not particular to Colonel Blotto games; the only criteria for its use is that decision spaces are discrete and utilities can be reasonably approximated by a normal distribution.

What it means to solve a model quickly is subjective, and for the purposes of this chapter this meant being able to solve a model on a commercial laptop running Python in one day. Using this criterion the analyis stopped after $n = 5$, but with a more powerful machine and/or willingness to wait for solutions larger problems could be solved. The main result of this chapter is that the OCBA-based algorithm offers enormous computational savings for ARA models and presents a framework for exploring the computational feasibility of increasingly larger models. A natural next step is to develop methods to identify promising regions of very large decision spaces, within which Algorithm 3.1 can be applied. Three general approaches come to mind for exploring this problem: assess all possible decisions using a low-fidelity model, a la Xu et al. (2016), and then explore in detail regions where the low-fidelity model suggested expected utility was high; partition the decision space and draw samples of decisions from each, which are then explored in detail to decide whether the region as a whole ought to be explored, a la W. Chen et al. (2014); and calibrating a distribution for the expected utility of a pair of decision maker and adversary decisions, $g(d,a,s) = u_D(d,a,s) \cdot p(s|d,a)$, via a Markov Chain Monte Carlo simulation, a la Bielza, Müller, and Insua (1999), and then analyzing the marginal of the distribution with respect to d to determine regions where the expected utility is high. All these approaches will have to deal with the unique challenge of ARA models, which is that the adversary's utility function and beliefs are stochastic. In this chapter this was overcome by utilizing a nested OCBA algorithm.

Another natural extension would be to incorporate a richer description of uncertainty regarding the adversary's behavior. Considering that the distribution of the adversary's behavior, $P_D(a|d)$, was estimated by solving his problem using statistical selection, despite the accuracy of the algorithm seen herein it's still subject to sampling error. By using simple point estimates for each $P_D(a_i|d)$, $a_i \in F_A$, sampling error could lead the decision maker to assume the adversary is using a suboptimal strategy, which in turn leads her to overestimate her utility from each strategy and may even suggest a false optimum. While it's reasonable to base decisions around Algorithm 3.1 because, by Theorem 3.1, when analyzing the adversary's problem it maximizes the probability of predicting his true optimum, risk-averse decision makers may instead wish to characterize the uncertainty in each $P_D(a_i|d)$ and employ a maximin strategy that maximizes utility subject to the worst (reasonable) case for the values of $P_D(a_i|d)$, even at the expense of lower expected utility. Such an approach might involve estimating confidence intervals for each $P_D(a_i|d)$ and using these as uncertainty sets in a distributionally robust optimization. Distributional robustness in ARA models was studied in McLay, Rothschild, and Guikema (2012), but in their analysis uncertainty sets were taken as given rather than estimated analytically. Intuitively, heuristically defined uncertainty sets such as $P_D(a_i|d) \in \{\hat{P}_D(a_i|d) \pm 5\%\}$, where $\hat{P}_D(a_i|d)$ are the point estimates used in this chapter, could be employed, but ideally future research will formalize rigorous methodologies to determine uncertainty sets.

Lastly, while Algorithm 3.1 was seen to outperform an alternative which is the only algorithm previously employed to a decently sized ARA problem, further alternatives should continue to be explored as computational ARA is a very young field. For example, the alternative classes of statistical selection techniques mentioned in Section 3.2 could be employed in the ARA context. This will allow the algorithm used here to be compared to other sophisticated methodologies to gain a better understanding of its value in solving this young class of game theoretic models.

Appendix 3.A. Closed Form Solutions for $\boldsymbol{\psi_D(d,a)}$ and $\boldsymbol{\Psi_A(d,a)}$

It's shown here that $\psi_D(d,a)$ and $\Psi_A(d,a)$ can be solved analytically using the distributions and utility functions specified in Section 3.4.1.

$$\psi_A(d,a) = \int_S U_A(d,a,s,r)p(s|d,a)ds \tag{3.10}$$

$$= \int_{h_n}^{h_n+.1} \int_{h_{n-1}}^{h_{n-1}+.1} \dots \int_{h_2}^{h_2+.1} \int_{h_1}^{h_1+.1} \frac{\sum_{i=1}^{n} r_i\left(1-e^{-4.6\left(s_i-C_{H,i}-.05\right)}\right)}{n} \cdot 10^n \, ds_1 \dots ds_n$$

(since $S_i|d,a \sim Uni(h_i, h_i+.1)$ for all i)

$$= \int_{h_n}^{h_n+.1} \dots \int_{h_2}^{h_2+.1} \left[\int_{h_1}^{h_1+.1} \frac{r_1\left(1-e^{-4.6\left(s_1-C_{H,1}-.05\right)}\right)}{n} \cdot 10\, ds_1 \right] \cdot$$

$$\frac{\sum_{i\neq 1} r_i\left(1-e^{-4.6\left(s_i-C_{H,i}-.05\right)}\right)}{n} \cdot 10^{n-1} ds_2 \dots ds_n.$$

Focusing just on the term inside the "[ ]:"

$$\int_{h_1}^{h_1+.1} \frac{r_1\left(1-e^{-4.6\left(s_1-C_{H,1}-.05\right)}\right)}{n} \cdot 10\, ds_1 = \frac{r_1}{n} \cdot \left(1 - 10 \cdot e^{4.6\left(C_{H,1}+.05\right)} \left[-\frac{1}{4.6} e^{-4.6 s_1}\right]_{h_1}^{h_1+.1}\right) \tag{3.11}$$

$$= \frac{r_1}{n} \cdot \left(1 + \frac{10}{4.6} \cdot e^{4.6\left(C_{H,1}+.05\right)} \cdot \left(e^{-4.6(h_1+.1)} - e^{-4.6 h_1}\right)\right)$$

$$= \frac{r_1}{n} \cdot \left(1 + \frac{10}{4.6} \cdot e^{4.6\left(C_{H,1}+.05\right)} \cdot (e^{-.46} - 1) \cdot e^{-4.6 h_1}\right)$$

$$\coloneqq I_1^A.$$

From here, it's easy to show that repeatedly integrating over the elements of S yields:

$$\psi_A(d, a) = \sum_{i=1}^{n} I_i{}^A. \tag{3.12}$$

At the time $\psi_A(d, a)$ must be evaluated, d and a are assumed known and hence h_i is known for all i. r_i is also assumed known at this point, and everything $C_{H,i}$ is a given constant. Hence, $\sum_{i=1}^{n} I_i^A$ can be calculated by direct computation. The defender's utility function has an almost identical form as the attacker's, and it's easy to show that $\psi_D(d, a) = \sum_{i=1}^{n} I_i{}^D$, where:

$$I_i^D \coloneqq -\frac{v_i}{n} \cdot \left(1 + \frac{10}{4.6} \cdot e^{4.6\left(C_{H,1}+.05\right)} \cdot (e^{-.46} - 1) \cdot e^{-4.6 h_1}\right). \tag{3.13}$$

Appendix 3.B. Formulation of $\boldsymbol{a^*(d) = argmax_a \psi_A(d, a)}$ as a Quadratic Integer Program

At node A of Figure 3.1.b, the attacker's problem is to solve:

$$\max_{a} z = \sum_{i=1}^{n} I_i{}^A \tag{3.14}$$

$$s.t.$$

$$\sum_{i=1}^{n} a_i = 1$$

$$a \in \{0,\ .1,\ .2,\ \dots,\ .9,\ 1\}.$$

Simply plugging in the formula for h_i given in Section 3.4.1 into (3.11) yields:

$$I_i^A = \frac{r_i}{n} \cdot \left(1 + \frac{10}{4.6} \cdot e^{4.6\left(C_{H,i}+.05\right)} \cdot \left(e^{-.46} - 1\right) \cdot e^{-4.6h_i}\right) \tag{3.15}$$

$$= A_{1,i} + A_{2,i} \cdot e^{-4.6h_i},$$

where $A_{1,i} \coloneqq \frac{r_i}{n}$ and $A_{2,i} \coloneqq A_{1,i} \cdot \frac{10}{4.6} \cdot e^{4.6\left(C_{H,i}+.05\right)} \cdot \left(e^{-.46} - 1\right)$.

$$= A_{1,i} + A_{2,i} \cdot e^{-4.6 \cdot -\frac{2}{4.6} \cdot \log\left(\frac{c_A^i a_i + 1}{c_D^i d_i + 1}\right) - 4.6 C_H^i}$$

$$= A_{1,i} + A_{2,i} \cdot e^{-4.6 C_H^i} \cdot \left[e^{\log\left(\frac{c_A^i a_i + 1}{c_D^i d_i + 1}\right)}\right]^2$$

$$= A_{1,i} + A_{2,i} \cdot e^{-4.6 C_H^i} \cdot \left(c_D^i d_i + 1\right)^{-2} \left(c_A^i a_i + 1\right)^2,$$

which is quadratic in a_i.

Therefore, model (3.14) is a quadratic program with a discrete feasible region. Making the transformation $a_i \rightarrow 10a_i$ and adjusting the objective function accordingly gives a small quadratic integer program which can be solved using off-the-shelf solvers.

Appendix 3.C. Stopping Criterion for Each Trial of Algorithm 3.1

In steps 15 and 18 of Algorithm 3.1, the following stopping criteria is used based on a weighted sum of percentage changes in strategies across iterations. The weights applied to each strategy are the number of samples allocated in the most recent iteration, normalized to sum to 1: $w_i = N_i / \sum_i N_i$. This places the vast majority of the weight only on those strategies reasonably thought to be the optimum. Using these weights, a weighted average percentage change between the current and each of the last four iterations is calculated, where percentage changes are defined with respect to the difference between the best and the worst expected utility at the current iteration, $\delta \coloneqq max_i \alpha_i - min_i \alpha_i$, where as in Section 3.4.3, α_i represents the sample expected utility for strategy i. This choice was made because dividing by expected utility itself causes unduly high percentage changes when α_i is near 0. Adding a subscript for the iteration number and assuming the current iteration is t, the weighted percentage change between the current iteration and iteration t' is: $q_{t,t'} \coloneqq \sum_i w_{t,i} \cdot \left|\alpha_{t,i} - \alpha_{t',i}\right| / \delta_t$. The stopping criteria is that $q_{t,t'} \leq .05$ for $t' = t-1, t-2, t-3$, and $t-4$, indicating no significant changes over the last four iterations.

Appendix 3.D. Proof of Lemma 3.1 and Theorem 3.2

3.D.1. Lemma 3.1

Wilson (1927) gives $1-\alpha$ one-sided confidence intervals for p_i of:

$$p_i \geq \frac{N - f_i + \frac{z^2}{2}}{N + z^2} - \frac{z}{N + z^2} \sqrt{\frac{f_i (N - f_i)}{N^3} + \frac{z^2}{4}}, \text{ where } z \coloneqq \Phi^{-1}(1-\alpha). \qquad (3.16)$$

By definition if $p_b > .5$, then E_b is the most likely outcome, and therefore if $p_b > .5$ with $1-\alpha$ confidence, then E_b is the most likely outcome with at least $1-\alpha$ confidence. The qualifier "at least" $1-\alpha$ confident is used, and this is a sufficient but not necessary condition, since it could be that $p_i < .5$ for all i.

3.D.2. Theorem 3.2

Part 1 of the theorem is proved by showing that: (1a) $2f_b + n_0 - 1 < N(f_b)$; and (1b) $N(f_b) \leq 2f_b + n_0 + 1$. Using Wilson's confidence intervals, part 1a is equivalent to showing:

$$\frac{f_b+n_0-1+z^2/2}{2f_b+n_0-1+z^2} - \frac{z}{2f_b+n_0-1+z^2}\sqrt{\frac{f_b(f_b+n_0-1)}{(2f_b+n_0-1)^3}+\frac{z^2}{4}} < .5 \leftrightarrow \tag{3.17}$$

$$\frac{n_0-1}{2} - z\sqrt{\frac{f_b(f_b+n_0-1)}{(2f_b+n_0-1)^3}+\frac{z^2}{4}} < 0 \leftrightarrow$$

$$\frac{f_b(f_b+n_0-1)}{(2f_b+n_0-1)^3} > \left(\frac{n_0-1}{2z}\right)^2 - \frac{z^2}{4}.$$

Define $g_c(x) = \frac{x(x+c)}{(2x+c)^3}$, which will be used again when proving parts 1b and 2 of the theorem. If $x, c \geq 0$, then $g(x) \geq 0$, and since $f_b, n_0 - 1 \geq 0$, if it's shown that $0 > \left(\frac{n_0-1}{2z}\right)^2 - \frac{z^2}{4}$, then part 1a will have been proved. Note that:

$$\left(\frac{n_0-1}{2z}\right)^2 - \frac{z^2}{4} = \frac{1}{4}\left(\frac{n_0-z^2-1}{z}\right)\left(\frac{n_0+z^2-1}{z}\right). \tag{3.18}$$

Clearly, $\frac{n_0+z^2-1}{z} > 0$. Further, $n_0 - z^2 < 1$ and therefore $\frac{n_0-z^2-1}{z} < 0$. It follows that $\frac{1}{4}\left(\frac{n_0-z^2-1}{z}\right)\left(\frac{n_0+z^2-1}{z}\right) = \left(\frac{n_0-1}{2z}\right)^2 - \frac{z^2}{4} < 0$, and therefore part 1a has been proved.

Similarly to part 1a, part 1b is true if $g_{n_0+1}(f_b) \leq \left(\frac{n_0+1}{2z}\right)^2 - \frac{z^2}{4}$ for all f_b. For any value of $c > 0$, calculus shows that $g_c(x)$ attains its maximum value at $x = \frac{c}{2}(\sqrt{3}-1)$, and $g_c\left(\frac{c}{2}(\sqrt{3}-1)\right) = \frac{\sqrt{3}}{18c}$. Therefore, to prove part 1b it's sufficient to show $\frac{\sqrt{3}}{18(n_0+1)} \leq \left(\frac{n_0+1}{2z}\right)^2 - \frac{z^2}{4}$. Observe:

$$\frac{\sqrt{3}}{18(n_0+1)} \leq \left(\frac{n_0+1}{2z}\right)^2 - \frac{z^2}{4} \leftrightarrow$$

$$\frac{2\sqrt{3}}{9} \leq (n_0+1)\left(\frac{n_0+1}{z} - z\right)\left(\frac{n_0+1}{z} + z\right) \leftrightarrow$$

$$\frac{2\sqrt{3}}{9} \leq (n_0+1)\left(\frac{n_0-z^2+1}{z}\right)\left(\frac{n_0+z^2+1}{z}\right) \leftrightarrow$$

$$\frac{2\sqrt{3}}{9} \leq \left(\frac{n_0+1}{z}\right)(n_0-z^2+1)\left(\frac{n_0+z^2+1}{z}\right). \tag{3.19}$$

Noting each term in the product on the righthand side of (3.19) is greater than 1 and that $\frac{2\sqrt{3}}{9} = .3849 < 1$, it follows that (3.19) is true, and therefore 1b is true.

Part 2 of the theorem is a natural extension of part 1. It's already known $N(f_b) > 2f_b + n_0 - 1$ by part 1a. To show $N(f_b) = 2f_b + n_0$, following the same reasoning as part 1b a sufficient condition is that $\frac{\sqrt{3}}{18n_0} \leq \left(\frac{n_0}{2z}\right)^2 - \frac{z^2}{4}$. The reason this is not a necessary condition is that the point where $g_{n_0}(f_b)$ attains its maximum, $f_{b,max} = \frac{n_0}{2}(\sqrt{3}-1)$, may not be an integer, and the number of failures by definition is an integer. A necessary and sufficient condition for $N(f_b) = 2f_b + n_0$ is that $\max\{g_{n_0}(f_{b-}), g_{n_0}(f_{b+})\} \leq$

$\left(\frac{n_0}{2z}\right)^2 - \frac{z^2}{4}$, where $f_{b-} := \left\lfloor\frac{\sqrt{3}}{18n_0}\right\rfloor$ and $f_{b+} := \left\lceil\frac{\sqrt{3}}{18n_0}\right\rceil$; this is true because $g_c(x)$ is unimodal for $x, c \geq 0$.

To prove part 3 it must show that: (3a) $P(N(f_b)|p_b) = P(!N(f_b - 1), f_b|p_b) \cdot p_b^2$; and (3b) $P(!N(f_b - j), i|p_b) = \sum_{k=k_{min}}^{k_{max}} P(!N(f_b - j - 1), k|p_b) \cdot p_b^{2-i+k} \cdot (1 - p_b)^{i-k} \cdot \binom{2}{i-k}$. To see that 3a is true, note $N(f_b) - N(f_b - 1) = 2f_b + n_0 - (2(f_b - 1) + n_0) = 2$, and therefore if f_b failures have already been observed through the first $N(f_b - 1)$ trials the next two must both be successes, and hence $P(!N(f_b - 1), f_b|p_b)$ is multiplied by p_b^2. If anything less than f_b failures had been observed through the first $N(f_b - 1)$ trials, then the algorithm would have terminated because, by assumption, only $N(f_b - 1)$ trials are required to terminate the algorithm when $f_b - 1$ failures have been observed, and thus the equality in 3a holds. To prove part 3b, first note that, given i failures occurred through the first $N(f_b - j)$ trials, the minimum number of failures through $N(f_b - j - 1)$ is the larger of $i - 2$ and the least number of failures that prevents the algorithm from terminating after $f_b - j - 1$ trials, or $f_b - j$. Hence, define $k_{min} = \max\{f_b - j, i - 2\}$. Also note the maximum failures through $N(f_b - j - 1)$ trials is the smaller of $N(f_b - j - 1) = 2(f_b - j - 1) + n_0$ and the number of failures given to have occurred through $f_b - j$ trials, or i, and hence $k_{max} := \min\{2(f_b - j - 1) + n_0, i\}$. Now 3b holds as a matter of definition, since there are 2 trials between $N(f_b - j - 1) + 1$ and $N(f_b - j)$, inclusive, the number of failures occurring between these trials is $i - k$, and there are $\binom{2}{i-k}$ ways for $i - k$ failures to occur.

Appendix 3.E. Robustness of Algorithm 3.1's Results

In Section 3.5 it was seen that for each value of n, a single run of the algorithm produced the true optimum in each case. Also recall that when $n = 2$ the algorithm terminated after 3 trials (and hence no trials produced a false optimum), when $n = 3$, 7 trials were required (and hence 2 failures occurred), when $n = 4$, 5 trials were needed (hence, 1 failure), and when $n = 5$ the algorithm required 3 trials (0 failures). To gain further understanding as to the algorithm's ability to produce the true optimum and the required number of trials to terminate, it was reran several times for the case when $n = 4$. Together with the results from Section 3.5 it was ran a total of 18 times for $n = 4$ using the model parameters in Section 3.5, which amounted to 102 total trials. In all 18 runs the true optimum was identified, but the number of trials required for termination could vary significantly. The required number was concentrated towards the low end, with a mode of just 3 (the minimum possible) and an average of 5.7, but the maximum required was 13, which occurred only once. Empirically, in over 102 trials the true optimum was selected 76.47% of the time. Also of interest is the average time required for each trial; this was 52 minutes, and hence in Section 3.5.3 the algorithm was benchmarked against a greedy search that took approximately $5.7 \cdot 52/60 = 4.92$ hours to terminate.

To add further credence to the algorithm, it was reran 10 times for three alternative parameter sets when $n = 4$, listed below. The required number of trials and time to termination are given in Table 3.2. Table 3.2 also shows the true utility of the strategy selected, given as a percentage of the true optimum; the majority of the time the algorithm selects the true optimum, but when it doesn't the true utility of the strategy selected is at a minimum 95.60% of the true optimum. Lastly, using only the samples generated (i.e. *not* using knowledge of the true optimum) Table 3.2 lists the lower bounds on the probability of correctly selecting the optimal strategy (APCS). As noted in Section 3.5.2, APCS can be a poor metric when multiple strategies provide nearly equivalent utilities. To provide decision makers with a sensible measurement that assumes ambivalence between strategies that are virtually equivalent, the following modified version of equation (3.6) is defined:

$$APCS(x) \coloneqq 1 - \sum_{i \in I} \Phi\left(\frac{\hat{u}_D(d_i) - \hat{u}_D(d_i)}{\sqrt{se_b^2 + se_i^2}}\right), \tag{3.20}$$

where se_i is the standard error of strategy i, $I \coloneqq \left\{i \mid \frac{\hat{u}_D(d_i) - \hat{u}_D(d_w)}{\hat{u}_D(d_b) - \hat{u}_D(d_w)} < x\right\}$, d_w is the strategy with the lowest sample expected utility, and as before d_b is that with the highest sample expected utility. Equation (3.20) effectively considers any strategy that's within $100 \cdot x$ percent of the highest sample expected utility to be an acceptable strategy choice. $APCS(x)$ was computed for $x \in \{99\%, 98\%, 97\%, 96\%, 95\%\}$ and the results are reported in Table 3.2.

Original parameters

$C_H = [.4,\ .35,\ .4,\ .4]$, $c_A = [-.4984, -.4984, -.5529, -.6015]$,

$c_D = [-.4984, -.4373, -.4373, -.5529]$, $v = [1.3,\ .8,\ 1.25,\ .7]$, and

$R = [R_1, R_2, R_3, R_4]$, where $R_1 \sim triangular(.8,\ 1,\ 1.5)$,

$R_2 \sim triangular(.5,\ .8,\ 2.5)$, $R_3 \sim triangular(1,\ 1.5,\ 3.5)$, and $R_4 \sim triangular(.3,\ .7,\ 1.1)$.

Parameter set 2: mirroring scenario. In this scenario, the modes in the triangle distributions for R equal v, and the upper and lower bounds are $\pm$.3 from the modes.

$C_H = [.4,\ .35,\ .4,\ .4]$, $c_A = [-.4984, -.4984, -.5529, -.6015]$,

$c_D = [-.4982, -.4373, -.4373, -.5529]$, $v = [1.3,\ .8,\ 1.25,\ .7]$, and

$R = [R_1, R_2, R_3, R_4]$, where $R_1 \sim triangular(1,\ 1.3,\ 1.6)$,

$R_2 \sim triangular(.5,\ .8,\ 1.1)$, $R_3 \sim triangular(.95,\ 1.25,\ 1.55)$, and

$R_4 \sim triangular(.4,\ .7,\ 1)$.

Parameter set 3: low incentive scenario. The status quo (C_H) is relatively good for the decision maker at all but her least desirable battlefield (battlefield 4).

$C_H = [.33, .33, .33, .4]$, $c_A = [-.5730, -.5210, -.6194, -.6015]$,

$c_D = [-.4108, -.4108, -.3390, -.5529]$, $v = [1.3,\ .8,\ 1.25,\ .7]$, and

$R = [R_1, R_2, R_3, R_4]$, where $R_1 \sim triangular(.8,\ 1,\ 1.5)$,

$R_2 \sim triangular(.5,\ .8,\ 2.5)$, $R_3 \sim triangular(1,\ 1.5,\ 3.5)$, and

$R_4 \sim triangular(.3,\ .7,\ 1.1)$.

Parameter set 4: randomized scenario. These parameters bear no relation to the original parameter set.

$C_H = [.39,\ .35,\ .37,\ .34]$, $c_A = [-.5827, -.5210, -.5529, -.5730]$, $c_D = [-.5631, -.3540, -.4242, -.4982]$, $v = [2.3,\ 1.7,\ 3.4,\ 2.1]$, and $R = [R_1, R_2, R_3, R_4]$, where $R_1 \sim triangular(2.8,\ 3,\ 3.1)$, $R_2 \sim triangular(2.5,\ 3.4,\ 3.5)$, $R_3 \sim triangular(.8,\ 1.1,\ 1.9)$, and $R_4 \sim triangular(2.9,\ 3.2,\ 3.5)$.

Table 3.2. Required trials, times, and lower bounds on PCS.

Parameter Set	Run	Trials	Time (hours)	% Optimum	APCS	APCS(99%)	APCS(98%)	APCS(97%)	APCS(96%)	APCS(95%)
Original	1	5	4.6209	100.00%	94.5781%	94.5781%	94.5781%	94.5781%	98.8322%	98.8322%
Original	2	3	2.3290	100.00%	86.4522%	86.4522%	86.4522%	86.4522%	86.4522%	88.6974%
Original	3	5	4.1744	100.00%	88.9501%	88.9501%	96.6762%	96.6762%	96.6762%	96.6762%
Original	4	3	2.5943	100.00%	65.9286%	65.9286%	65.9286%	65.9286%	67.3257%	72.0249%
Original	5	3	3.1451	100.00%	91.9086%	91.9086%	91.9086%	96.3086%	96.3086%	96.3086%
Original	6	3	2.6621	100.00%	48.1133%	88.0593%	88.0593%	88.0593%	88.0593%	88.0593%
Original	7	9	7.4437	100.00%	98.6345%	98.6345%	98.6345%	98.6345%	99.9408%	99.9408%
Original	8	5	4.3114	100.00%	98.3496%	98.3496%	98.3496%	98.3496%	98.3496%	99.8269%
Original	9	7	6.0717	100.00%	82.9616%	82.9616%	82.9616%	99.8365%	99.8365%	99.8365%
Original	10	5	4.0994	100.00%	93.3289%	93.3289%	93.3289%	93.3289%	94.1222%	98.8892%
Original	11	3	2.7186	100.00%	51.7583%	51.7583%	51.7583%	70.9374%	78.3933%	80.5176%
Original	12	11	9.5211	100.00%	96.5826%	96.5826%	99.6875%	99.9922%	99.9922%	99.9922%
Original	13	3	2.4231	100.00%	72.9906%	72.9906%	87.7637%	87.7637%	87.7637%	87.7637%
Original	14	5	4.7192	100.00%	94.6204%	94.6204%	94.6204%	94.6204%	97.3383%	97.4005%
Original	15	13	11.0425	100.00%	98.1009%	98.1009%	98.1009%	99.8232%	99.9682%	99.9977%
Original	16	9	7.6797	100.00%	99.0674%	99.0674%	99.0674%	99.9343%	99.9343%	99.9590%
Original	17	7	6.2415	100.00%	98.5667%	98.5667%	98.5667%	98.5667%	99.8249%	99.8844%
Original	18	3	2.6933	100.00%	58.1801%	58.1801%	58.1801%	70.5259%	70.5259%	72.6036%
2	1	3	2.4675	100.00%	99.9025%	99.9025%	99.9025%	99.9025%	99.9025%	99.9025%
2	2	3	2.5464	100.00%	99.9719%	99.9719%	99.9719%	99.9719%	99.9719%	99.9719%
2	3	3	2.7767	100.00%	99.9982%	99.9982%	99.9982%	99.9982%	99.9982%	99.9982%
2	4	3	2.5408	100.00%	99.9464%	99.9464%	99.9464%	99.9464%	99.9464%	99.9464%
2	5	3	2.6114	100.00%	99.9463%	99.9463%	99.9463%	99.9463%	99.9463%	99.9463%
2	6	3	2.8472	100.00%	99.9823%	99.9823%	99.9823%	99.9823%	99.9823%	99.9823%
2	7	3	2.7889	100.00%	99.9977%	99.9977%	99.9977%	99.9977%	99.9977%	99.9977%
2	8	5	4.1136	100.00%	99.2218%	99.2218%	99.2218%	99.2218%	99.2218%	99.9732%
2	9	3	2.4742	100.00%	99.9795%	99.9795%	99.9795%	99.9795%	99.9795%	99.9795%
2	10	3	2.4822	100.00%	99.9357%	99.9357%	99.9357%	99.9357%	99.9357%	99.9357%
3	1	5	4.2419	98.85%	20.7777%	62.9565%	97.8979%	97.8979%	97.8979%	97.8979%
3	2	3	2.5672	99.82%	47.6974%	47.6974%	69.1484%	79.8305%	81.9716%	94.8479%
3	3	3	2.3906	99.75%	93.2065%	93.2065%	96.2494%	96.9947%	99.1739%
3	4	11	8.8308	99.75%	74.1395%	74.1395%	98.9298%	99.7562%	99.9864%	99.9995%
3	5	7	5.9297	99.75%	20.7633%	99.4819%	99.4819%	99.4819%	99.9517%	99.9517%
3	6	3	2.5475	100.00%	-2.3521%	65.6943%	86.5776%	86.5776%	90.7887%	96.7025%
3	7	5	4.1925	100.00%	73.1479%	73.1479%	94.4568%	98.8500%	98.8500%	99.3940%
3	8	3	2.4447	100.00%	28.7584%	66.4094%	92.8074%	92.8074%	93.8089%	95.2626%
3	9	3	2.3878	100.00%	35.7291%	35.7291%	74.9243%	86.2074%	86.2074%	91.3297%
3	10	3	2.6775	100.00%	-10.8486%	68.6894%	88.7759%	88.7759%	94.3021%	94.3021%
4	1	5	4.2408	99.84%	5.9631%	68.2740%	76.6379%	91.0776%	98.2393%	99.9059%
4	2	3	2.3983	95.60%	36.1192%	36.1192%	58.7280%	72.0863%	90.5764%	96.8366%
4	3	3	2.3967	99.84%	-49.0188%	62.2326%	62.2326%	72.3289%	88.1562%	88.1562%
4	4	3	2.8375	100.00%	-64.3456%	27.3455%	66.3419%	81.2045%	91.5331%	91.5331%
4	5	3	2.4619	100.00%	-65.9713%	14.3660%	14.3660%	54.6571%	93.9504%	93.9504%
4	6	7	6.5617	96.72%	48.0206%	74.3938%	94.3866%	99.5608%	99.7396%	99.7656%
4	7	9	7.9350	100.00%	16.8710%	73.2499%	98.9508%	98.9508%	98.9508%	99.9212%
4	8	9	7.5425	99.84%	32.1717%	89.2021%	89.2021%	99.2128%	99.6805%	99.9842%
4	9	15	12.9708	100.00%	25.3987%	80.5595%	99.6900%	99.6900%	99.9772%	99.9952%
4	10	9	8.1714	100.00%	95.0478%	95.0478%	98.2804%	98.2804%	99.9421%	99.9962%

To benchmark results against the greedy search algorithm of Section 3.5.3, $APCS(x)$ was calculated for the 5 values of x in Table 3.2 when the original parameters were used. These evaluated to $APCS(99\%) = APCS(98\%) = APCS(97\%) = -82.00\%$, and $APCS(96\%) = APCS(95\%) = -19.76\%$. Recall that APCS proper for the greedy algorithm was -82.00%. In order to generate a bound greater than 95%, x must be set to 81% or lower, effectively taking the (unreasonable) position of being ambivalent between utilities that are within 81% of the true optimum.

4 ANALYZING THE SOUTH CHINA SEA FISHING DISPUTE AS A COMPLEX GAME: EFFICIENT SAMPLE ALLOCATION VIA A RESPONSE SURFACE METHODOLOGY

Note. This chapter has been submitted for publication as a stand-alone article in *Computational Economics*.

China, Vietnam, Malaysia, Brunei, the Philippines, and Taiwan all have overlapping territorial claims in the South China Sea (SCS). Among the issues making the SCS a valuable natural resource is fish, where confrontations between the nations' coast guards and fishermen have become routine. This chapter studies this dispute in the game-theoretic context by formulating a model for allocating coast guard patrol craft across multiple fisheries as a way to impose costs on rival fishermen. The game is complex because the SCS has many distinct fisheries, the biomasses of which behave as nonlinear processes (and generally, more-than-quadratic). The most complicating factor in the model, however, is uncertainty in how the players behave. As critics of game theory have noted, humans aren't necessarily so "rational" that the standard game-theoretic solution concepts have meaning in the real world. More generally, if the underlying model is off slightly, the results could in theory differ greatly from real world outcomes. To address each of these issues a novel evaluation criterion for games is introduced. It's akin to robust analysis and, not surprisingly, makes an already computationally intensive game much more difficult. Thus, a response surface methodology is used to analyze the game in a feasible amount of time.

4.1. Introduction

The South China Sea (SCS) is one of the most profitable bodies of water for commercial fishing and has become the source of territorial disputes between its bordering countries. Traditional United Nations defined exclusive economic zones (EEZs), which extend 200 nautical miles beyond a state's coastline, are problematic in a congested maritime environment such as the SCS where China, Vietnam, Malaysia, Brunei, the Philippines, and Taiwan have all made territorial claims which overlap with those of at least one other state. This situation has led to more than mere diplomatic jostling; real, physical clashes have occurred between fishermen and the coast guards and navies of rival nations, and in the most extreme cases have led to the loss of life. The economic implications are also enormous as encroachments by one nation's fishermen into waters claimed by another are a daily occurrence. The parties in question have taken actions to improve their standing in the dispute. On the one hand, some countries have equipped their fishermen with martial assets as a form of "maritime militia" to mitigate the deterrent effect of rival maritime patrols. On the other, all nations other than China appear to be willing to cooperate to resolve the dispute via the multilateral fisheries management organization the Southeast Asian Fisheries Development Center (SEAFDEC). The U.S. has also taken an interest in the dispute on account of its rivalry with China and the broader implications for control of the SCS.

Given the SCS is vast and composed of multiple fisheries, it's natural to model this dispute as a many-variate game of strategic resource allocation. Each fishery can be modeled using a bionomic model and the additional costs imposed by maritime patrols can be informed by recent empirical research on the subject. There are multiple modeling challenges, however. First, even the simplest bionomic models won't lead to analytic solutions when multiple fisheries are in dispute, so sampling methods will be required. Second, decently sized problems become computationally intensive and thus efficiency of sampling will be paramount. While not commonly used in the game theory literature, this first pair of problems can actually be addressed straightforwardly using well-established techniques in response surface methodologies (RSM). A third problem is a bit more nuanced and requires ingenuity.

Traditionally, a game-theoretic model would predict the players' behavior using a Nash equilibrium, or near-Nash equilibrium, but this can hide key risks when deviations from such model-prescribed behavior leads to large deviations in realized utilities. Deviations in behavior could result for several reasons:

1. Foremost, the underlying model is almost surely not a perfect representation of the player's utility functions, and deviations should thus be expected.
2. Even if the model were perfect, if the players being modeled don't possess the tools to find true optimal responses they'll rather "satisfice" by finding a good solution that might be suboptimal. This creates a host of possible strategies where the players might achieve stability, rather than a single true (near) Nash equilibrium.

This chapter introduces a novel metric to assess the risks stemming from such uncertainty in behavior, referred to as "behavioral uncertainty" for the remainder of the chapter. For a variable upper-bound on the proximity to a Nash equilibrium, maximum and minimum attainable utilities will be found, thus capturing the range of possible outcomes for a given tolerance for how far players deviate from model-directed behavior. This added layer of complexity heightens the need for efficient sample selection, and the sampling algorithm developed shows superior performance against two simple benchmarks.

The remainder of the chapter is organized as follows. Section 4.2 reviews the literature on the SCS dispute, fishery games, concepts related to behavioral uncertainty, and response surface methodologies in general. Section 4.3 presents the general model for the optimal response of a player, which is assumed to be arbitrarily complex and solvable only via a derivative-free, computationally expensive optimizer. Section 4.4 presents the sampling algorithm to analyze the game with motivating examples for the need to assess risks stemming from uncertainty in players behavior. Section 4.5 provides examples to show the algorithm's performance, and Section 4.6 concludes and comments on future research.

4.2. Literature Review

The South China Sea accounts for approximately one-tenth of all fish caught globally, making it one of the most economically important natural resources on the planet ("South China Sea Threatened by 'a Series of Catastrophes'" 2019). More broadly, the SCS is economically and politically important due to its strategic implications for shipping, energy resources, and its potential military importance during war (Buszynski 2012; Yoshihara 2012). Even under a more mild set of conditions, geography alone would likely lead to disputed claims of ownership as traditionally defined EEZs overlap. In fact, each China, Vietnam, Malaysia, Brunei, the Philippines, and Taiwan have made claims that overlap with those of at least one other nation (Stearns 2012). The competitive environment in the SCS has contributed to the rapid growth in China's coast guard to the point where it's larger than that of its neighbors combined (Erickson 2018), and to the expansion of state-sponsored maritime militias, which are commercial fishermen armed with small-arms and water cannons intended to offset the deterrent effects of patrols (Erickson and Kennedy 2016; Zhang and Bateman 2017). Instances of fishermen clashing with coast guard patrols at sea are common; often, rival coast guards will confront one another in response to a distress signal sent by a fisherman (China Power Team 2020). While this is a multilateral dispute, it's often thought of through a China-versus-the-rest lens. This isn't entirely inappropriate, especially when viewed in terms of the implications for fisheries management and the emergence of SEAFDEC, an intergovernmental fisheries management organization including the major SCS nations other than China. This chapter will take this view and present a two-player game; note, however, the sampling methodology developed is not contingent on the game having only two players.

From a modeling perceptive, bionomic models of fisheries represent a well-established field and these models have been applied to relatively simple games to draw general insights. Introductory texts on bionomic fisheries models and game theory in fisheries management are Clark (2006) and Grønbæk et al. (2020), respectively. Grønbæk et al. (2020) covers noncooperative and cooperative games, two- and many-layer games, asymmetries, and discounted payoffs. Advanced complicating factors such as multiple interacting species (Fischer and Mirman 1996), coalition formation in many-player games (Long and Flaaten 2011), and uncertainty in stock levels and growth rates (Miller and Nkuiya 2016) have also been analyzed using games. The most general conclusion drawn is that, as in other games modeling common-pool resources, fish become overexploited in a competitive environment.

Optimal maritime patrol allocation has also been analyzed using game theory, though not through the lens of territorial disputes. Rather, these models have taken fishing rights as given and determined how to allocate a single state's patrols to combat illegal fishing (Fang, Stone, and Tambe 2015; M. Brown, Haskell, and Tambe 2014). A related strand of research has conducted empirical analysis on the extent to which patrols (and other factors) can deter illegal fishermen (Petrossian 2015). Crucially, empirical analysis has deliberately left out data on the SCS because it's not clear the same relationships will govern undisputed and disputed EEZs. Consider, for instance, the ability of maritime militias and interacting patrols to offset the effects of a single state's patrols. The true nature of the effect of patrols in the SCS is therefore an open question;

this chapter presents a simple model that in theory could be estimated with data, but that empirical research is left for follow-on research.

What was defined as behavioral uncertainty in Section 4.1 has been studied in a variety of ways. For instance, quantal response models assume players are more likely to play strategies with higher utilities, though they aren't required to play the true optimal strategy with probability one (McKelvey and Palfrey 1995; Nguyen et al. 2013). In recent years adversarial risk analysis (ARA) has emerged as a method incorporating uncertainty in the model parameters to derive probabilistic distributions of players' behaviors; this derivation generally occurs by drawing Monte Carlo samples for the model parameters and then fitting an empirical distribution to realized behavior (Rios Insua, Rios, and Banks 2009; Banks, Rios, and Rios Insua 2016). The approach used in this chapter to address behavioral uncertainty is both distribution- and parametric-free, making it akin to robust analysis. Robust games, generally defined as assessing the worst-case scenario for an adversary's behavior, have been addressed from a theoretical perspective in Aghassi and Bertsimas (2006), Kardes (2005), and Crespi, Radi, and Rocca (2020). Robustness has been applied to realistic, though small, games in Brown, Haskell, and Tambe (2014) and McLay, Rothschild, and Guikema (2012).

All the methodologies mentioned above to account for behavioral uncertainty either impose restrictive parametric forms on the game, do not scale well to large problems, or both, and hence the focus of this chapter on developing a response surface methodology to analyze large, complex games. RSM approximates a function that's time consuming to compute, such as a player's optimal response, using an instantaneously computable function which is calibrated to a small sample of realizations of the original function. RSM has long been popular for approximating difficult optimizations and a taxonomy of methods was presented by Jones (2001). This taxonomy decomposes RSM into interpolating and non-interpolating methods, and one-stage and two-stage methods where the latter generates subsequent sample points following an initial sampling of the decision space. Polynomial regressions and other parametric models have traditionally been used as response surfaces but black-box models such as boosted regression trees and artificial neural networks can be incorporated as well (Hastie, Tibshirani, and Friedman 2009). Two-stage methods involve optimizing the response surface to pick subsequent sample points; this is feasible when using the aforementioned black-box models given cutting-edge global optimizers and increased computing power (Xu et al. 2015).

4.3. The Fishing Dispute Game with Multiple Fisheries

Two players, Blue and Red, are competing over k fisheries and must allocate a fixed amount of maritime patrols to each. This is akin to the SCS where the largest state, China, is largely in competition with a coalition of the others. Section 4.3.1 introduces the fisheries model governing rent extracted from a fishery, and Section 4.3.2 states the optimal response problem of one player given the maritime patrol allocation of the other.

4.3.1. A Fisheries Model with Costs Imposed by Patrols

The biomass, x_i, of each fishery $i = 1{:}\,k$ is modeled by the following model of growth and decay:

$$\frac{dx_i}{dt} = r_i x_i \left(1 - \frac{x_i}{Z_i}\right)^{\alpha_i} - q_i x_i^{\gamma_i} F_i. \tag{4.1}$$

r_i is the natural growth rate of fish, Z_i is the fishery's maximum carrying capacity, F_i is the total amount of fishing in fishery i, and q_i is the "catchability coefficient." When $\alpha_i = \gamma_i = 1$, this is the common Gordon-Schaeffer model for fisheries. $\alpha_i \neq 1$ allows the dependency between the natural growth rate and biomass to be nonlinear, and $\gamma_i \neq 1$ allows one to model "patchy" populations. Populations that are more patchy ($\gamma_i < 1$) remain relatively easy to catch when population declines because the remaining stocks stick together, while non-patchy populations become harder to catch as population declines because a smaller population is spread over the same geographic space. The long-term biomass, given F_i, is found by setting (4.1) equal to 0, which can be computationally solved virtually instantaneously. This long-term biomass will be denoted $\tilde{x}_i$.

Due to overlapping claims in the SCS, each player independently determines how many fishing quotas to issue for each fishery. These choices will be made with the aim of maximizing the sum of resource rents from each fishery, which for Blue is defined as:

$$\pi_B(P_B, P_R) = \sum_{i=1:k} \pi_{B,i} \tag{4.2}$$

$$\pi_{B,i} = \left(p_i q_i \tilde{x}_i^{\gamma_i} - \psi_{B,i}\right) F_{B,i}, \tag{4.3}$$

where $P_B = \left[P_{B,1}, P_{B,2}, \dots, P_{B,k}\right]$ and $P_R = \left[P_{R,1}, P_{R,2}, \dots, P_{R,k}\right]$ hold Blue and Red's patrol allocations for each fishery, $F_{B,i}$ is the amount of quotas Blue allocated to fishery i, p_i is the price fetched for one metric ton of fish from fishery i, and $\psi_{B,i}$ is the total cost of fishing for a Blue fishermen in fishery i.

$\psi_{B,i}$ of course accounts for standard operating and opportunity costs, but because the fisheries are in dispute costs imposed by patrols are also influential. The following linear model is used for costs:

$$\psi_{B,i} = c_B + \max\{0,\ \beta_{BR} P_{R,i} - \beta_{BB} P_{B,i}\}. \tag{4.4}$$

The term β_{BR} quantifies the effect Red patrols impose on Blue fishermen, and β_{BB} quantifies the ability of Blue patrols to offset that effect. The net effect of patrols on costs obviously cannot be negative, so the $\max\{\cdot\}$ function is used to keep costs at or above Blue's level of operating and opportunity costs, c_B. For simplicity it's been assumed a patrol allocated to fishery i does not have the range to influence other fisheries. While in practice this may not be accurate, this assumption doesn't affect the methodological approach of Section 4.4. Equations (4.2), (4.3), and (4.4) can analogously be defined for Red.

Once the patrol allocations for each player are given, it's possible to derive Nash equilibrium fishing levels in each fishery if $\alpha_i = \gamma_i = 1$. In the general case, equilibrium levels can be modeled as a function of (P_B, P_R) using a truncated polynomial model. To see this, consider the following:

Result 4.1. Equilibrium fishing levels, given patrols.

1. Define $F^*_{B,i}(F_{R,i}, P_{B,i}, P_{R,i})$ and $F^*_{R,i}(F_{B,i}, P_{B,i}, P_{R,i})$ as the optimal fishing levels for Blue and Red, respectively, given the other's fishing levels and each's patrol allocation. Define $x_i(F_{B,i}, F_{R,i})$ as the biomass given fishing levels $F_{B,i}$ and $F_{R,i}$.
2. $F^*_{B,i}(F_{R,i}, P_{B,i}, P_{R,i}) = 0$ iff $p_i q_i x_i\left(0, F^*_{R,i}(0, P_{B,i}, P_{R,i})\right)^{\gamma_i} - \psi_{B,i} \leq 0$. Otherwise, $F^*_{B,i}(F_{R,i}, P_{B,i}, P_{R,i})$ can be modeled using a non-computationally expensive design of experiment (DOE) over $(F_{R,i}, P_{B,i}, P_{R,i})$; a polynomial regression is seen to fit sampled data with accuracy $R^2 \approx .99$. The same holds for modeling $F^*_{R,i}(F_{B,i}, P_{B,i}, P_{R,i})$.
3. Given the near perfect models for optimal response fishing levels in point 2, a non-computationally expensive DOE can be performed over (P_B, P_R) to find equilibrium fishing levels as a function of patrols. For each sample, first check if $F_{B,i}$ or $F_{R,i}$ are provably 0. If neither are, standard nonlinear methods can find an equilibrium in less than one second, using the polynomial models for optimal responses. Fitting a further polynomial for equilibrium fishing levels as a function of patrols has accuracy $R^2 \approx .99$.

The sampling methods described in Result 4.1 don't add a material amount of time to the broader problem of optimally allocating patrol vessels across many fisheries, as described in Section 4.4. Denote the equilibrium fishing levels modeled via Result 4.1 as $\tilde{F}_{B,i}$ and $\tilde{F}_{R,i}$. Recalling that $\tilde{x}_i$ is a function of $F_i = F_{B,i} + F_{R,i}$, and $\tilde{F}_{B,i}$ a function of

$P_{B,i}$ and $P_{R,i}$, Blue's rent (and analogously, Red's) can be written strictly in terms of patrols:

$$\pi_B(P_B, P_R) = \sum_i \left(p_i q_i^{\gamma_i} \tilde{x}_i - c_B - \max\{0,\ \beta_{BR} P_{R,i} - \beta_{BB} P_{B,i}\}\right) \tilde{F}_{B,i}. \qquad (4.5)$$

4.3.2. Optimal Response Function for patrols

The optimal response functions for Blue and Red are the following:

$$P_B^*(P_R) \coloneqq \underset{P_B}{\operatorname{argmax}}\, \pi_B(P_B, P_R)\ \ s.t. \sum_{i=1:k} P_{B,i} = P_B^{tot} \qquad (4.6)$$

$$P_R^*(P_B) \coloneqq \underset{P_R}{\operatorname{argmax}}\, \pi_R(P_B, P_R)\ \ s.t. \sum_{i=1:k} P_{R,i} = P_R^{tot}, \qquad (4.7)$$

where P_B^{tot} and P_R^{tot} are the total number of patrols at Blue and Red's disposal. These functions are costly optimization problems, where the costs stem from needing to compute $\tilde{x}$ computationally, the need to model equilibrium fishing levels via a high-degree polynomial, and the potentially high value of k. Specifically, when $k = 10$ these optimizations require an average of over three minutes to solve on a standard laptop, and when $k = 60$ can take over 90 minutes. Whether the purpose of analysis is to seek a pure Nash equilibrium or some measure of proximity to equilibrium, because no analytical solution to (4.6) and (4.7) exists these functions will need to be computed for several values of (P_B, P_R). An efficient sampling technique is paramount and is the subject of the next section.

4.4. Response Surface Methodology for Analyzing the Fishing Dispute Game

4.4.1. Measuring Proximity to the Model-Defined Equilibrium

Consider the following metric for the proximity of player strategies to a true Nash equilibrium:

$$e(P_B, P_R) \coloneqq \sqrt{\frac{\left(\pi_B\left(P_B^*(P_R), P_R\right) - \pi_B(P_B, P_R)\right)^2 + \left(\pi_R\left(P_B, P_R^*(P_B)\right) - \pi_R(P_B, P_R)\right)^2}{2}}. \tag{4.8}$$

This metric simply measures the root mean squared deviation of the players' realized utilities from the optimal attainable utilities from a unilateral move. Clearly, at a Nash equilibrium $e = 0$.

Recall the points made in Section 4.1 for why seeking a pure Nash equilibrium, or a single near-Nash equilibrium if one does not exist, does not suffice. First, the underlying model is at best an approximation of actual player behaviors. A Nash equilibrium might be a sensible forecast of behavior, but any number of strategies may also be reasonable. Second, players may satisfice rather than finding true optimal responses. In each of these two cases, it seems sensible that pairs of strategies where $e(P_B, P_R)$ is small are more likely. For this reason the following pair of optimizations is defined which will quantify how variable realized Blue rents are as a function of how large e is allowed to be:

$$\max_{P_B, P_R} \pi_B(P_B, P_R) \ \ s.t. \, e(P_B, P_R) \le d \tag{4.9}$$

$$\min_{P_B, P_R} \pi_B(P_B, P_R) \ \ s.t. \, e(P_B, P_R) \le d. \tag{4.10}$$

(4.9) and (4.10) give the largest and smallest possible Blue rents, respectively, such that e remains less than some pre-specified level. Solving these optimizations for various values of d will capture how the risks stemming from behavioral uncertainty change as behavior is allowed to deviate further from model-specified optima. As motivating examples, consider Figures 4.1.a and 4.1.b, which were produced using the RSM described in Section 4.4.2. Figure 4.1.a plots the solutions to (4.9) and (4.10) for increasing values of d using an example where $k = 3$ (line 16 in Table 4.1). In this case, a pure Nash equilibrium exists as seen by the convergence of the upper- and lower-bounds at $d = 0$. However, modest values of d cause the bounds to diverge significantly. When d = \$2 million (only 2.11% of the rent Blue realizes at the Nash equilibrium), the upper-bound on π_B is 16.11% higher than the lower-bound. Figure 4.1.b illustrates an example where no Nash equilibrium exists (line 1 in Table 4.1). Here, the smallest value of d for which the RSM found two valid pairs of strategies is \$3.25 million, which is quite small relative to the lower-bound on Blue rent at this value of d, \$319 million. Nevertheless, the range of possible rents is quite large: the upper-bound is 28.67% higher than the lower-bound.

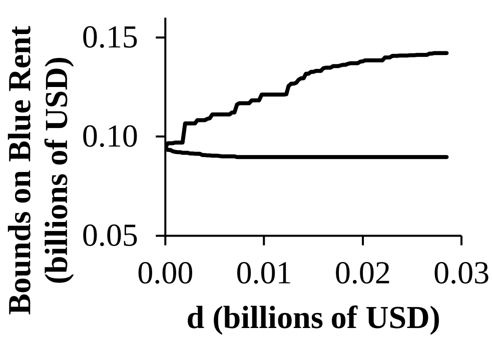

Fig. 4.1.a. Existence of a true Nash equilibrium

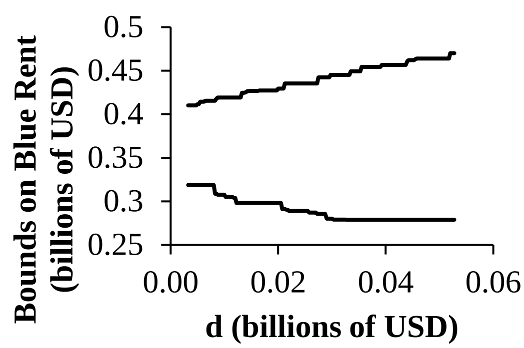

Fig. 4.1.b. No Nash equilibrium exists

Figure 4.1. Bounds on π_B for increasing values of d in (4.9) and (4.10).

4.4.2. Response Surface Methodology for Optimizations (4.9) and (4.10)

Given the lack of an analytical solution to (4.9) and (4.10) and the computation times noted in Section 4.4.1, a brute force search across the strategy space is infeasible for decently sized problems (e.g. $k = 10$). This is true even if one only wants to solve (4.9) and (4.10) for a single value of d, much less a continuum of values as in Figure 4.1. In light of this burden, a sampling method is required to analyze (4.9) and (4.10) for multiple values of d. The essence of any such method can be encapsulated in Algorithm 4.0:

Algorithm 4.**0**

1. Draw samples of (P_B, P_R). For each sample:
   a. Compute $\pi_B(P_B, P_R)$ and $\pi_R(P_B, P_R)$.
   b. Compute $\pi_B(P_B^*(P_R), P_R)$ by solving (4.6). Similarly compute $\pi_R\big(P_B, P_R^*(P_B)\big)$ via (4.7).
   c. Compute $e(P_B, P_R)$ using the results of steps 1.a and 1.b.
2. Having drawn these samples, simply observe the maximum and minimum values of $\pi_B(P_B, P_R)$ such that $e(P_B, P_R) < d$, for varying values of d.

The algorithm used in this chapter takes an iterative approach to step 1. An initial set of samples is drawn, a response surface is then fit to the initial sampling, and subsequent samples are drawn that are likely to be the most informative based on information obtained via the response surface. This is captured formally in Algorithm 4.1:

Algorithm 4.1

1. Generate initial samples as follows:

    a. Generate n_{LHS} samples of (P_B, P_R) using a modified Latin hypercube sample. Specifically, generate points by allowing $P_{B,i}, P_{R,i} \in [0,1]$ for all fisheries i, and then scale each sample so $\sum_i P_{B,i} = P_B^{tot}$ and $\sum_i P_{R,i} = P_R^{tot}$.

    b. Generate n_{uni} samples of (P_B, P_R) based on unilateral responses computed via a fast heuristic. Specifically, define $P_B'(P_R)$ and $P_R'(P_B)$ as easily computable functions that give good responses for Blue to Red, and Red to Blue, respectively, though not necessarily optimal responses. Then: (i) generate a random sample of (P_B, P_R); (ii) randomly select a player, WLOG Blue, to move first; (iii) the sample point to be used to compute e is $\left(P_B'(P_R), P_R'\left(P_B'(P_R)\right)\right)$. That is, Blue responds to the randomly generated P_R via a fast heuristic, and Red in turn responds via a fast heuristic. This sampling approach could conceivably provide a more representative sample of strategies employed near an equilibrium.

    c. For each sample point from steps 1.a and 1.b, use a global optimizer to find $P_B^*(P_R)$ and $P_R^*(P_B)$. Then, compute $e(P_B, P_R)$.

2. Fit a response surface for $e(P_B, P_R)$ using currently drawn samples, denoted $\hat{e}(P_B, P_R)$. If the fit is adequate, go to step 3. Otherwise, generate one more

sample using the method in step 1.a, one more using the method of step 1.b, and repeat step 2.

3. Define a discrete set $\vec{d} = [d_1, d_2, \dots, d_N]$, where $0 < d_i < d_{i+1}$. For $d_i = 1,2, \dots, N$, perform the following:
    a. Generate a random value of d using the uniform distribution $Uni(d_{i-1}, d_i)$. By definition consider $d_0 = 0$.
    b. Use a global optimizer to solve the response surface analogue of (4.9):
    $$\max_{P_B, P_R} \pi_B(P_B, P_R) \;\; s.t. \hat{e}(P_B, P_R) < d. \tag{4.11}$$
    Denoting the solution $(P_B^{*,max}, P_R^{*,max})$, compute $e(P_B^{*,max}, P_R^{*,max})$, and add this to the set of samples. Recalibrate $\hat{e}(P_B, P_R)$.
    c. Repeat step 3.b, using a minimization rather than a maximization in (4.11).
4. Repeat step 3 until a predefined computational budget is exhausted, then observe the maximum and minimum values of $\pi_B(P_B, P_R)$ such that $e(P_B, P_R) < d$, for varying values of d.

The next section provides several examples implementing Algorithm 4.1. In all cases, the following details apply:

- The computational budget was set at 300 samples, regardless of k.
- Examples where $k \leq 10$ used $n_{LHS} = n_{uni} = 2k$ in steps 1.a-b. A few larger examples were performed with $k = 30$ and $k = 60$. To ensure the bulk of

samples were still generated via the RSM, n_{LHS} and n_{opt} were set to 20 when $k = 30$, and 40 when $k = 60$.

- The fast heuristics in step 1.b are a COBYLA optimization (Powell 1998) with a random starting location and no restarts.
- The response surface in step 3 is a boosted regression tree (Hastie, Tibshirani, and Friedman 2009). In step 2, the number of trees to use and their depths is chosen via cross validation, and the criterion determining whether the fit is adequate is a cross-validated R^2 exceeding 0.5. To save time, cross validation is only performed in step 2 (not in steps 3.b and 3.c when recalibrating the surface).
- In step 3, $N = 5$, the points $0, d_1, d_2, \dots, d_5$ are evenly spaced, and d_3 is set equal to the maximum value of $e(P_B, P_R)$ found in the initial sampling.
- COBYLA with 20 random restarts is used in steps 3.b and 3.c.

For comparison, 300 samples were drawn while solely using the sampling techniques described in each steps 1.a and 1.b. The next section will demonstrate the value added of the RSM-based algorithm over these benchmark approaches.

4.5. Examples

4.5.1. Baseline Parameter Values

For all examples of the fishing dispute game considered, the parameters Z_i, r_i, q_i, α_i, and γ_i will not vary across fisheries and will simply be denoted Z, r, q, α, and γ. The prices of fish found in each fishery, measured in billions of dollars per metric ton, will vary from 1 to 3 in each example according to the following formula: $p_i = 1 + (i - 1) \cdot \frac{2}{k-1}$ (i.e. linearly increasing prices from \$1 to \$3 billion). Several other parameters are held fixed across examples and are provided here: $r = .4$, $q = .0002$, $c_B = c_R = .12$ (measured in billions of dollars per unit of fishing), $P_B^{tot} = 9k$, $P_R^{tot} = 15k$. The values for total patrols were chosen so Red would have a comparable number of patrols to the Chinese Coast Guard when $k = 60$, the largest example performed (Erickson 2018).

4.5.2. Example: Bounds on Rent, Effectiveness of Increased Patrols, and Implications for Negotiations

In this example, $k = 60$, $\alpha = \gamma = 1$, $\beta_{BR} = 2 \times 10^{-6}$, $\beta_{RB} = 1.33 \times 10^{-6}$, and $\beta_{BB} = \beta_{RR} = 8.33 \times 10^{-7}$. The coefficients on patrols were chosen such that $\beta_{RB} < \beta_{BR}$ to reflect the fact China has invested more in maritime militias than other states, and thus ought to be less effected by adversarial patrols. Figure 4.2 plots the upper- and lower-bounds of π_B as estimated via Algorithm 4.1 (labeled "RSM"). Also plotted are the bounds estimated using only the modified LHS technique of step 1.a ("LHS"), and those estimated via the unilateral response sampling technique of step 1.b ("uni"). The benefit of Algorithm 4.1 is apparent from the figure. For instance, when d = \$0.1 billion (which is small relative to π_B), the lower- and upper-bounds for π_B found by Algorithm 4.1 are \$1.82 billion and \$2.49 billion, while the alternative methods estimate bounds of just \$1.98 billion to \$2.06 billion (for LHS) and \$1.90 billion to \$2.05 billion (the unilateral response method).

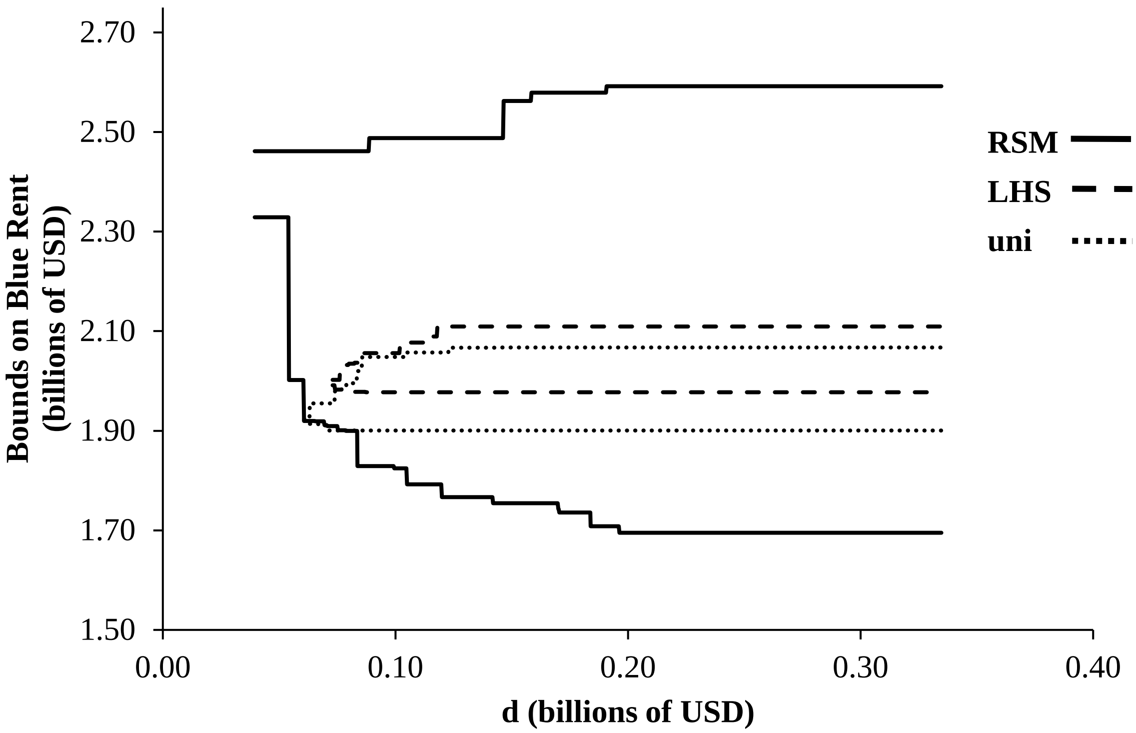

Figure 4.2. Bounds on π_B for three sampling methods.

Several more examples are summarized in Table 4.1, but before proceeding two policy-relevant questions are addressed which leverage the information obtained via Algorithm 4.1. First, Blue may want to assess the potential gain from increasing her overall number of patrol craft. Figure 4.3 plots both the upper- and lower-bounds on π_B for the current example's parameters, and those resulting from doubling Blue's patrols. Considering that the space between the bounds are the possible realized rents, the figure reveals that most outcomes that are possible from doubling patrols are also possible from not making further investments. Further, the lower-bound obtained from doubling patrols is strictly less than the upper-bound from not doubling. This is a conservative, though valid, way to assess the potential benefit of increased patrols, as behavioral uncertainty presumes the underlying model for Red's behavior is unknown, and hence a change in P_B^{tot} has an unknown effect. Even with a perfect model, satisficing may lead to unpredictable effects. Note this is only a static analysis of the effect of doubling Blue's patrols: in reality one should assess whether Red will respond with increased patrols of his own, or by taking actions to effect β_{RB} or β_{RR}. This isn't undertaken in this chapter because much more goes into the decision to invest in patrols and militias than their effects on fishing, such as search-and-rescue operations, port security, and trafficking in persons.

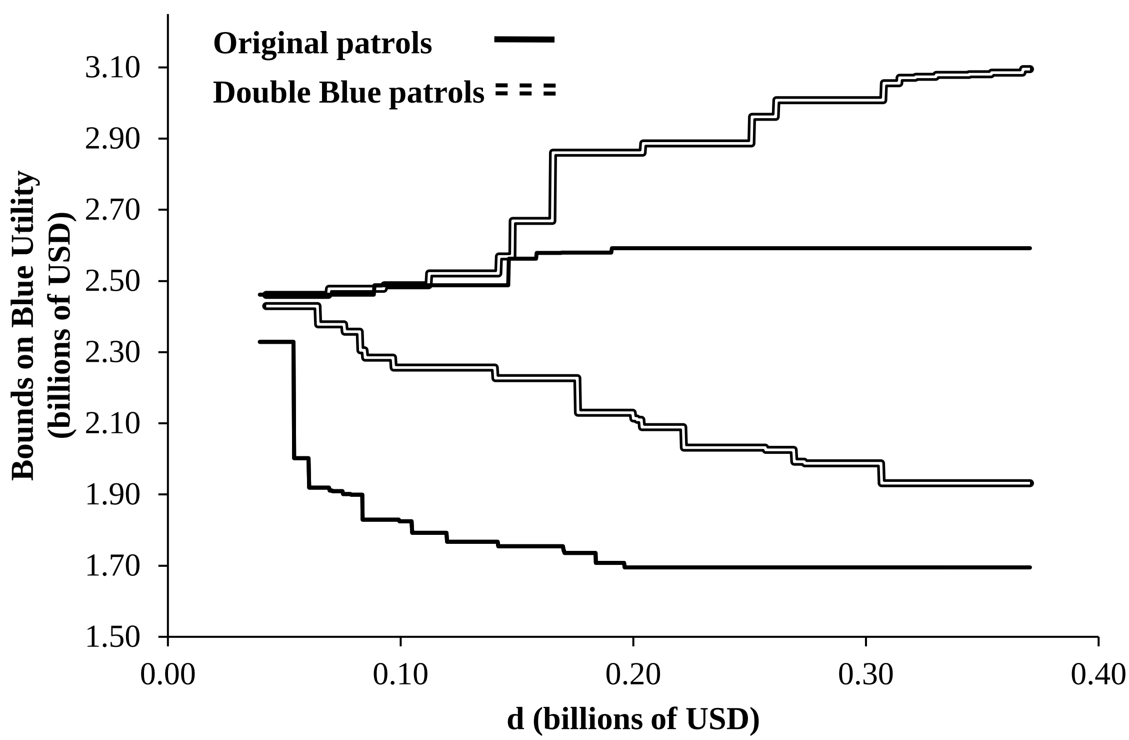

Figure 4.3. Effect of doubling Blue patrols.

The second application explored concerns negotiations. Blue and Red can each increase their rents by coming to a negotiated settlement for who is allowed to fish where, and in what quantities. A common model for negotiations is the Nash bargaining problem (Nash 1953):

$$\max_{F_B,F_R}(\pi_B - \bar{\pi}_B)(\pi_R - \bar{\pi}_R), \tag{4.12}$$

where $F_B = [F_{B,1}, F_{B,2}, \dots, F_{B,k}]$, $F_R = [F_{R,1}, F_{R,2}, \dots, F_{R,k}]$, and $\bar{\pi}_B$ and $\bar{\pi}_R$ are the "threat values" Blue and Red earn in the noncooperative disputed fisheries game. The rents π_B and π_R no longer depend on patrols as all agreed upon quotas are legal. As a brief illustration of the effects of behavioral uncertainty on negotiations, (4.12) was solved 100 times for randomly selected values of $\bar{\pi}_B \in [1.79, 2.48]$ billion USD and $\bar{\pi}_R \in [2.88, 3.22]$ billion USD (the latter was informed by running Algorithm 4.1 to obtain bounds on Red's realized rents). The results are unsurprisingly sizable, with π_B ranging from \$2.32 billion to \$2.77 billion, and π_R from \$3.22 billion to \$3.68 billion.

4.5.3. Additional Examples

Table 4.1 and Figure 4.4 provide several additional examples confirming the results of Section 4.5.1. The model parameters k, α, γ, Z, β_{BR}, β_{RB}, β_{BB}, and β_{RR} were all varied. Table 4.1 shows how much larger the upper-bound on π_B is relative to the lower-bound for various values of d. To contextualize the results and keep examples on a similar scale, the metric $\bar{\pi}_B$ is defined as the average Blue rent among all samples found via Algorithm 4.1 where e is in the bottom 5$^{\text{th}}$ percentile. For increasing values of $d/\bar{\pi}_B$, the percentage change between the bounds, $\left(\pi_B^{UB}(d) - \pi_B^{LB}(d)\right)/\pi_B^{LB}(d)$, is listed. Figure 4.4 provides a scatter plot of these percentage changes found via Algorithm 4.1

versus those found via the benchmark approaches (for each example and value of d, only the benchmark which found the greater percentage change is plotted). In the plot, different markers are used for different values of k.

The examples in Table 4.1 use non-Gordon Schaeffer parameter settings (that is, $\alpha \neq 1$ and $\gamma \neq 1$) and larger carrying capacities and smaller β coefficients than the example in Section 4.5.2. Values of α greater than 1 imply growth rates are largest when biomass is less than $Z/2$, and as seen in the table this has the effect of making behavioral uncertainty greater. In contrast, $\gamma \neq 1$ doesn't appear to be an influential parameter for behavioral uncertainty. Larger carrying capacities and smaller β coefficients (implying more profitable fishing and lesser effects of patrols, respectively), each decrease behavioral uncertainty. For all examples, the differences between lower- and upper-bounds remain large. A 45-degree line is plotted in Figure 4.4 to illustrate how much larger the difference in bounds found via Algorithm 4.1 is than that found by the best benchmark: markers above the line indicate Algorithm 4.1 identified a larger range of possible rents. With the exception of a single marker lying below the line when $k = 10$, the only ones not well above it correspond to $k = 3$. This indicates a naïve sampling method performs well for only the smallest problems. Once the strategy space reaches even a modest size, such as when $k = 5$, the benchmark methods miss key sections of the strategy space. The lone marker below the line can be explained by random sampling error, which is unavoidable in any sampling strategy.

Table 4.1

Additional examples for the range of possible Blue rents,

$\left(\boldsymbol{\pi_B^{UB}} - \boldsymbol{\pi_B^{LB}}\right)/\boldsymbol{\pi_B^{LB}}\ \boldsymbol{s.t.}\ \boldsymbol{e} < \boldsymbol{d} \cdot \bar{\pi}_B$

k	α	γ	β_{BR}	Z	d = 2.5%	d = 5%	d = 7.5%	d = 10%
10	1	1	2×10^{-6}	1	36%	44%	59%	64%
			1.3×10^{-6}	1	21%	24%	27%	27%
			2×10^{-6}	2	15%	17%	17%	17%
	1.5	1.25	2×10^{-6}	1	4%	27%	63%	90%
			1.3×10^{-6}	1	26%	44%	49%	51%
			2×10^{-6}	2	23%	25%	25%	25%
	1.5	0.75	2×10^{-6}	1	36%	45%	59%	65%
			1.3×10^{-6}	1	22%	27%	28%	28%
			2×10^{-6}	2	20%	20%	20%	20%
	0.5	1.25	2×10^{-6}	1	27%	30%	35%	35%
			1.3×10^{-6}	1	16%	16%	16%	16%
			2×10^{-6}	2	10%	10%	10%	10%
	0.5	0.75	2×10^{-6}	1	22%	25%	25%	25%
			1.3×10^{-6}	1	11%	11%	11%	11%
			2×10^{-6}	2	11%	11%	11%	11%
3	1	1	2×10^{-6}	1	5%	19%	23%	30%
			1.3×10^{-6}	1	11%	19%	22%	23%
			2×10^{-6}	2	8%	10%	10%	10%
5	1	1	2×10^{-6}	1	34%	44%	50%	54%
			1.3×10^{-6}	1	16%	21%	25%	25%
			2×10^{-6}	2	13%	13%	13%	13%
30	1	1	2×10^{-6}	1	30%	43%	50%	58%
			1.3×10^{-6}	1	22%	27%	27%	27%
			2×10^{-6}	2	16%	20%	20%	20%
60	1	1	2×10^{-6}	1	6%	36%	46%	53%
			1.3×10^{-6}	1	19%	25%	26%	29%
			2×10^{-6}	2	15%	19%	20%	20%

Note. In all examples, $\beta_{RB} = (2/3) \cdot \beta_{BR}$, and $\beta_{BB} = \beta_{RR} = (\beta_{BR} + \beta_{RB})/4$. Parameters not noted equal the baseline values provided in Section 4.5.1.

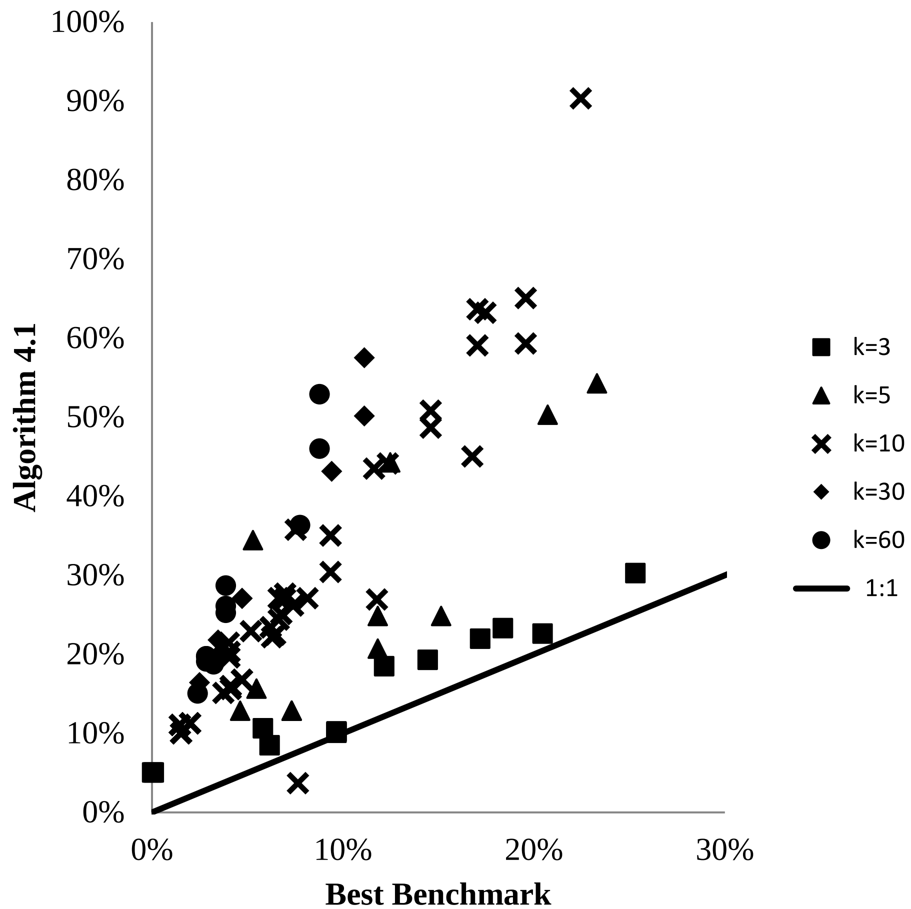

Figure 4.4. Percentage differences of bounds estimated via Algorithm 4.1 vs. benchmarks.

4.6. Conclusion and Future Work

This chapter has introduced a complex model for the fishing dispute in the South China Sea, where the strategic allocation of maritime patrols for two players impacts realized rents. The analysis wasn't limited to seeking an idealized, single point (near) Nash equilibrium, but rather sought to characterize the range of possible realized rents as a functions of proximity to equilibrium. This modelling framework was deemed necessary to account for what was termed behavioral uncertainty.

The formulated model is complex in the sense that an analytical solution does not exist and can only be analyzed by drawing computationally expensive samples. In additional to highly nonlinear fishery models and a budget constraint on patrols, much of the computational difficulty was driven by the presence of multiple fisheries in the SCS, and the chapter successfully analyzed examples with up to 60 fisheries. This was done by developing a sampling algorithm incorporating response surface methodologies, which was seen to vastly outperform two benchmarks.

A general finding is that behavioral uncertainty plays a large role in realized rents across the breadth of examples. In particular, when the root mean squared deviation of player rents from their optimal attainable rents (i.e. e) is just 2.5% of realized rents closest to a true equilibrium, the percentage increase from the lower- to upper-bound on rent can be as high as 36%. When the root mean squared deviation is relaxed to 10%, the difference in bounds can reach 90%. Refer to Table 4.1. An example was performed to illustrate how these results effect the decision to invest in additional patrols and the results of a negotiated agreement on fishing rights, and it was seen that behavioral uncertainty ought to affect both significantly.

This chapter didn't attempt to derive the traditional theoretical results found in the RSM literature, such as guarantees for convergence, based on the presumption that sufficiently complex problems are better modeled by complex response surfaces that won't lend themselves to such analytical results. In this case, a boosted regression tree was used but similarly complex surfaces are available. A strand of future research would be to perform extensive experiments on a wide variety of problems, possessing a variety of overarching features, and attempt to identify factors where certain types of response surfaces perform well. For example, a tree-based method performed well in this game of strategic resource allocation. This may persist into other games of strategic resource allocation, as optimal behavior may suddenly shift in a discontinuous way once one battlefield becomes dominated by one player.

One other point of future research is to operationalize the methodology developed here by engaging fisheries experts to truly model the SCS. This chapter was experimental and tested an RSM-based methodology across several instantiations of fishery parameters. It also assumed fisheries were dispersed enough such that each patrol vessel could be allocated to one and only one fishery at a time. A virtue of the methodology used in this chapter is no assumptions were imposed on the underlying model, hence this latter point requires no modification to the methodology. Another empirical point, mentioned earlier, is that research must still be performed to understand how patrols impose costs on fishermen in a disputed fishery; currently, the best research has only examined the effect of patrols on illegal fishermen under undisputed territorial rights. By accurately modeling the parameters governing the SCS, and the effects of patrols on costs, policy makers will have a valuable tool for answer critical questions surrounding this fisheries dispute.

5 FISHERIES MANAGEMENT IN CONGESTED WATERS: A GAME-THEORETIC ASSESSMENT OF THE EAST CHINA SEA

Note. This chapter is currently under peer-review as a stand-alone article with *Environmental and Resource Economics*.

The East China Sea (ECS) poses a different challenge than the patrol allocation problem of the South China Sea analyzed in the last chapter. In the ECS, fishing rights between China, South Korea, and Japan have essentially been agreed to. However, fishermen from one country encroach on the rights of another on a daily basis. Legal agreements bind the countries into punishing their own fishermen for such violations, but the continued, routine nature of encroachments suggest something deeper is going on. This chapter develops a game theoretic model where the nations' lax monitoring, controls, and surveillance (MCS) policies allow illegally caught fish to enter ports without repurcussion. Lax MCS thus gives nations a tacit way of illegally extracting rent from another's waters. Recognizing this, policy makers must negotiate not only on territorial rights in congested seas (as in Chapter 4), but also implement a mutually observable MCS policy. Without each, illegal fishing will persit and rents, it's seen, will be lower. Methodologically this scenario can be modeled as a binary program and an algorithm is developed to find a provable Nash equilibrium when each of two nations own one fishery. With multiple fisheries, the model becomes a nonconvex binary program, something not often seen in the game theory literature where linear and convex quadratic programs predominate. A modern global optimization technique (nested partitioning) is used in the multi-fishery case.

5.1. Introduction

The East China Sea (ECS) is a congested maritime environment in the sense that fishermen from multiple nations can access these waters with essentially equivalent operating costs. A journey from China to Japan through the ECS, for example, can span less than 350 nautical miles, while one from China to South Korea can span only 200 nautical miles. As a point of fact, fishermen from each of these three nations have historically fished throughout the ECS under less restrictive fisheries management regimes. Other congested maritime environments can of course be identified, such as the South China Sea (SCS) where disputes exist between all major SCS nations, but a distinguishing feature of the ECS is that legal agreements have been made on paper for who's allowed to fish where (Rosenberg 2005).

Despite such agreements between each China, South Korea, and Japan, illegal encroachments persist in high volumes (Hsiao 2020). Writing this off as the behavior of self-interested criminals is unsatisfactory, as state-led monitoring, controls, and surveillance (MCS) of fisheries present cost-effective ways to deter illegal fishing, and yet MCS has been rated poorly among ECS nations by the United Nations Food and Agricultural Organization (FAO) (Petrossian 2019). If these illegal encroachments were a true concern of the state, one would see more investment in MCS. This chapter provides a game-theoretic explanation for how rational state policy can allow illegal encroachments to persist.

The chapter analyzes a congested maritime environment using a game with two players, Blue and Red. A key modeling framework is that despite legally defined fishing waters, nations can tacitly allow illegal fishing by issuing excessive quotas in their own waters while letting economic principles determine whether each fisherman will fish in Blue waters, Red waters, or neither. Coupled with a standard fisheries model, this economic determination is influenced by asymmetries in costs. While the proximities of nations are assumed to render operating costs equivalent, asymmetries exist in the ECS on at least two accounts. Opportunity costs are the most obvious asymmetry, as income opportunities in China differ greatly from South Korea and Japan. The other source of asymmetry addressed is that arising from the deterrent effect of patrol crafts. Not only can nations have differences in the quantity and quality of patrols, which impose costs on illegal fishing, but a nation's fishermen also differ in their ability to resist patrols. The latter point has garnered significant attention in recent years with the rise of so-called maritime militias; that is, fishermen armed with martial assets (often with state sponsorship) such as water cannons and small arms that embolden them to fish in illegal waters. As will be shown, the low-cost player can extract rent from the high-cost player's waters by issuing excessive quotas; one strategy the high-cost player can use to combat this is to issuing excessive quotas herself, reducing the attractiveness of her waters to illegal fishermen and creating a vicious circle of overfishing.

As already alluded, fisheries management entails more than issuing the "right" amount of quotas. MCS measures are also essential to ensure fishermen are fishing only where allowed, catching the allowable species in the allowed quantities, and so on. For the purposes of this chapter, MCS refers to such deterrent practices other than maritime patrols and includes measures such as port inspections and vessel monitoring systems. It does not include sound practices for issuing quotas, which is treated as a separate decision variable. The poorly-rated MCS observed in the ECS can also be explained through strategic interaction in a congested environment. It's seen that even if MCS is a cost-effective deterrent of illegal fishing in the single-state context, once a second player is introduced the incentive to invest in MCS vanishes.

A critical component in the analysis is that each player can behave against the norms of agreed upon fishing rights while maintaining plausible deniability. For example, a country, say China, cannot explicitly authorize fishermen to fish in South Korean waters, as if this were discovered the political consequences would be unacceptable. This assertion is consistent with observed behavior as China has historically punished its fishermen caught in South Korean waters (Zhang 2016). Likewise, if it were discovered China were knowingly letting fish caught in South Korean waters in through its ports, the consequences would be unacceptable. It might be asked, then, what is to be done, if an agreement is already on paper but countries nonetheless tacitly violate it? To address this, the chapter assesses what enforceable policy on fishing quotas can be implemented to improve each player's utility. Such a policy involves an observable agreement on legal fishing levels, including legally allowed fishing by one player in the other's waters, as well as mutually observable MCS regimes.

Unsurprisingly, the low-cost player, either by virtue of low opportunity costs or strong patrols and maritime militias, has leverage in negotiating such an enforceable policy. However, it's found that in many situations only a reasonable amount of legally allowed fishing by the low-cost player in the high-cost player's waters can lead to a significant and balanced benefit for each, and thus there's likely a politically amenable solution to improve fisheries management in the ECS.

The remainder of the chapter is structured as follows. Section 5.2 reviews the literature on fisheries management in the ECS as it pertains to illegal fishing, game theoretic models of fisheries, and the deterrent effects of maritime law enforcement on illegal fishing. Section 5.3 describes a model for the ECS as a two-player game where each player owns a specified number of fisheries; this includes a subgame to determine where fishermen with quotas will fish. Section 5.4 derives analytical results for the case when each player owns a single fishery, and examples are presented for each the single- and multi-fishery cases. Section 5.5 concludes and comments on future work. Two appendices are provided with mathematical proofs and solution algorithms for the game.

5.2. Literature Review

Bilateral agreements delineating fishing rights in the ECS began to take form in the late 1990s, delineating regions that would fall under the jurisdiction of particular states as well as jointly managed areas. These agreements have also included provisions for legal fishing by one country in the waters of another (Rosenberg 2005; Hsiao 2020). However, these agreements are not fully cooperative in the sense that they haven't set limits on domestic fishing; that is, a bilateral agreement between China and South Korea may specify legal Chinese fishing levels in South Korean waters, but not legal Chinese levels in Chinese waters. This oversight has led to overfishing, and is perhaps the cause of current cutbacks on legally allowed foreign fishing (*Yonhap News* 2020). Despite these agreements, illegal encroachments persist in large numbers. An estimated 29,600 instances of Chinese fishing vessels entering South Korean waters occurred in the second half of 2014 alone. Further, and in support of the assertion states may only tacitly encourage illegal fishing but not outright allow it, Chinese officials have taken punitive action against its illegal fishermen when detained by South Korea and other countries (Zhang 2016). Analyzing these disputes has been complicated by the rise of maritime militias, the most well-studied of which is the Chinese Maritime Militia. These militias are fishermen who've been equipped with martial assets such as small arms and water cannons, generally through state-sponsorship, to oppose other nation's fishermen and patrol craft (Erickson and Kennedy 2016; Zhang and Bateman 2017; Perry 2020). In the most extreme cases, interactions between fishermen and maritime law enforcement has led to the loss of life on both sides (China Power Team 2020; Park 2020).

A key tool for addressing illegal encroachments, and illegal fishing in general, is monitoring, controls, and surveillance (MCS), which includes measures such as inspections at ports of entry, onboard observers to monitor fishing activities, and satellite systems to track vessels (Pitcher, Kalikoski, and Pramod 2006). The definition of MCS also often includes the issuance of quotas, but quotas issuance is treated separately in this chapter. Empirical studies have shown MCS to be an effective deterrent to illegal fishing, and while costs can vary widely irrespective of MCS quality its generally viewed as a cost-effective tool (Petrossian 2015; Mangin et al. 2018). Fisheries management experts have, nonetheless, rated MCS practices among ECS nations as below average (Petrossian 2019). Interestingly, studies have also found more generally that fisheries management tends to be weaker when nearby states are also not adhering to best practices, lending credence to the modeled results in this chapter for congested maritime environments (Borsky and Raschky 2011).

From a modeling perspective fisheries disputes have often been studied through game theory. In the well-studied problem of two or more players allocating quotas in the same fishery, models have shown overfishing will occur both when considering discounted payoffs and not. When accounting for asymmetric costs a player's utilities are seen to increase as another's costs increase (Grønbæk et al. 2020). While related to the findings of this chapter, the distinction is the model presented here considers multiple fisheries with legal definitions of who's allowed to fish in each, introducing a criminological element into the analysis. Models have accounted for additional complicating factors such as multiple interacting species (Fischer and Mirman 1996), coalition forming in many player games (Long and Flaaten 2011), and uncertainty in stock levels and growth (Miller and Nkuiya 2016), to name a few. Models have also been used to optimize patrol strategies to combat illegal fishing. These models discretize a nation's waters and consider two players, the state and illegal fishermen, who must determine how to deploy patrols and where to attempt illegal fishing, respectively. Factors such as heterogeneous illegal fishermen types, bounded rationality, and repeated play have all been incorporated (Fang, Stone, and Tambe 2015; M. Brown, Haskell, and Tambe 2014). What has not been considered is situations with multiple states competing over resources utilizing either state-sponsored or tacitly encouraged illegal fishing.

Regarding the economics of illegal fishing, researches have assessed the annual dollar amount of illegally caught fish to be between \$10-\$23.5 billion (Petrossian 2019). In addition to MCS, empirical studies have found the number of patrol craft per 100,000 square kilometers to be a significant explainer of illegal fishing (Petrossian 2015). Deterrence has also been analyzed at the agent-level, where studies have shown fishermen do indeed rationalize their illegal fishing by pointing out the risks and costs associated with being caught don't outweigh the benefits (King and Sutinen 2010; Kuperan and Sutinen 1998; Nielsen and Mathiesen 2003). Lastly, note a gap in the empirical literature exists in understanding the impact of maritime militias on the willingness to fish illegally; while maritime militias have been studied extensively, no empirical evidence exists to quantify their impact.

5.3. The Congested Environment Fishing Model

The scenario is modeled as a two-player game between Blue (B) and Red (R). In the general case, each player can have ownership of multiple fisheries. Each fishery is modeled using the Gordon-Schaeffer model with a logistic growth function and constant-effort harvesting rate (Clark 2006). The model for the j^{th} fishery of player $i \in \{B, R\}$, which will be referred throughout this chapter as fishery ij, is:

$$\frac{dx_{ij}}{dt} = r_{ij}x_{ij}\left(1 - \frac{x_{ij}}{Z_{ij}}\right) - x_{ij}q_{ij}\sum_{kl}F_{ij,kl}, \tag{5.1}$$

where x_{ij} is the current biomass in the fishery, Z_{ij} is the carrying capacity, r_{ij} the natural growth rate, and q_{ij} the catchability coefficient. $F_{ij,kl}$ is the level of fishing effort in the fishery by fishermen who are authorized to fish in fishery kl (the l^{th} fishery of player $k \in \{B, R\}$). In other words, the variable $F_{ij,kl}$ for $i \neq k$ or $j \neq l$ represents illegal fishing, whereas $F_{ij,ij}$ represents legal fishing. No migration between fisheries is assumed. In this chapter long-term biomass in each fishery is used to evaluate utilities. Given all values of $F_{ij,kl}$, the solution to the differential equation (5.1) is:

$$x_{ij} = Z_{ij}\left[1 - \frac{q_{ij}\sum_{kl}F_{ij,kl}}{r_{ij}}\right]. \tag{5.2}$$

To reiterate the key modeling framework, Blue and Red's national strategy merely determines how many fishing quotas are allocate in their respective fisheries, denoted F_{kl} for each fishery kl. It's assumed that without any explicit instruction from the state, individual fishermen with quotas decide where to fish based on economic principles. The variables $F_{ij,kl}$ will be determined via a subgame to be detailed momentarily.

An individual fishermen's revenue is a function of biomass. Denoting the price of fish farmed from fishery ij as p_{ij}, the revenue earned by a fisherman in this fishery is

$p_{ij}q_{ij}x_{ij}$. That fishermen's rent is determined by subtracting out costs, which obviously includes operational costs, but also asymmetric opportunity costs and costs imposed through the efficacy of maritime patrols to deter illegal fishing. It's assumed all fisheries are in close enough proximity such that operating costs are the same in each and for each player. In sum, the costs for a fisherman to fish in ij when he's legally authorized to fish in kl is defined as follows:

$$c_{ij,kl} = \begin{cases} c_k, & \text{if } i = k \text{ and } j = l \\ c_k + \beta_k P_i + \beta_m m_k, & \text{otherwise} \end{cases}, \tag{5.3}$$

where c_k is the combined operating and opportunity costs for fishermen from country k, P_i is the number of patrol craft employed per 100,000 square kilometers by country i, and m_k is the level of MCS used in country k. The coefficient β_m represents the deterrent effect of MCS on illegal fishing and is assumed constant for both players. β_k, in contrast, is the deterrent effect of patrols and is allowed to vary between players. This reflects differing levels of maritime militias among ECS nations. Considering these costs, the rent earned by a fisherman is $\pi_{ij,kl} = p_{ij}q_{ij}x_{ij} - c_{ij,kl}$.

Total utilities realized by Blue and Red, u_B and u_R, are the sum total of rents collected by their nation's fishermen, minus expenditures on MCS. MCS costs are modeled as a function of the level of MCS used; the flexible form $a_1 m_k^{a_2}$ is used, where a_1 and a_2 are constants. The level of patrols is treated as a given constant rather than a decision variable, as in practice many other considerations affect investment in patrol craft; for instance, port security, detection of human trafficking, and search and rescue operations, to name a few applications. The utility functions for Blue and Red are therefore:

$$u_B = \sum_{ij} \sum_{Bl} \pi_{ij,Bl} - a_1 m_B^{a_2} \tag{5.4}$$

$$u_R = \sum_{ij} \sum_{Rl} \pi_{ij,Rl} - a_1 m_R^{a_2}. \tag{5.5}$$

By solving a subgame for the values of $F_{ij,kl}$, these utility functions collapse into expressions of the players' overall decision variables (F_{Bl}, m_B, F_{Rl}, and m_R).

5.3.1. Subgame for Levels of Fishing in Each Fishery

A state-issued fishing quota gives fishermen the legal right to bring fish into the country of issue. States cannot explicitly direct fishermen to fish illegally; the decision on where to fish in a congested environment is therefore left to the fishermen, a choice which will be made based on economic principles. Elementally, the following must hold for a fishermen's choice to be economically rational: (i) a fishermen will use a quotas iff he can earn nonnegative rent; (ii) if he's fishing in fishery ij, he must be earning at least as much rent as he could earn in any other fishery $i'j'$, as otherwise he'd switch to the other. These subgame principles can be modeled by the below binary program, where the legally authorized quotas in each fishery, F_{Bl} and F_{Rl}, and the levels of MCS used, m_B and m_R, are given first-stage decision variables. The binary program finds values of the subgame variables (SGVs) $F_{ij,kl}$ such that points (i) and (ii) are satisfied for all fishermen.

$$\min_{z,\ F_{ij,kl},\ y_{ij,kl},\ y_{kl}} z \tag{5.6}$$

$s.t.$

Noninformative objective function (5.6a)

$z \geq 0$

Constraints on total fishing (5.6b)

$\sum_{ij} F_{ij,kl} < F_{kl}\ \forall\ l, k \in \{B, R\}$

Fishing only occurs where profitable, and is at least as profitable as any alternative fishery (5.6c)

$F_{ij,kl} \leq My_{ij,kl} \ \forall\ ij,\ kl$

$F_{ij,kl} \geq .0001y_{ij,kl} \ \forall\ ij,\ kl$

$\pi_{ij,kl} \geq -M\left(1 - y_{ij,kl}\right) \forall\ ij,\ kl$

$\pi_{ij,kl} \geq \pi_{i'j',kl} - M\left(1 - y_{ij,kl}\right) \forall\ ij,\ kl,\ k'l'$

No profitable quotas are left unused (5.6d)

$\sum_{ij} F_{ij,kl} + My_{kl} \geq F_{kl} \ \forall\ kl$

$\sum_{ij} F_{ij,kl} + .0001y_{kl} \leq F_{kl} \ \forall\ kl$

$\pi_{ij,kl} \leq M(1 - y_{kl}) \ \forall\ ij,\ kl$

Nonnegativity and binary constraints (5.6e)

$F_{ij,kl} \geq 0 \ \forall\ ij,\ kl$

$y_{ij,kl}, y_{kl} \in \{0,1\}.$

In the binary program (5.6), the parameter M is an arbitrarily large constant. The objective function is irrelevant as the purpose of the subgame is to find feasible values of the SGVs. When each player has only one fishery, unique subgame equilibriums can be proven to exist (see Theorem 5.1). Extensive random sampling indicates the guaranteed existence of a subgame equilibrium when each player has multiple fisheries, and while these subgame equilibria aren't guaranteed to be unique in general, they are unique at all equilibrium points for the overall game found in Section 5.4.2. An analytical formula for the subgame equilibrium can be found in the one-fishery case. These results are formalized in Theorems 5.1 and 5.2.

Theorem 5.1. Existence of a unique subgame equilibrium.

When Blue and Red each own only one fishery, any instantiation of the game's parameters and choice of the overall game's decision variables yields a unique subgame equilibrium.

Proof. See Appendix 5.A.

Theorem 5.2. Formulae for the unique subgame equilibrium via a partitioning of the parameter and decision space.

Again assume Blue and Red each own one fishery. The parameter and decision variable space can be partitioned, such that in each region of the partition analytical formulae exist for the subgame variables.

Proof. See Appendix 5.A.

5.4. Examples

5.4.1. One Fishery per Player

Because each player has only owns one fishery, subscripts $i1, k1$ are replaced with i, k, and $k1$ is replaced by k, for $i, k \in \{B, R\}$. This most basic instantiation of the model is sufficient to understand why nations overfish, and why illegal encroachments persist, in a congested environment. In addition to the analytical results in Theorems 5.1 and 5.2, when each player has only one fishery it can be proven $m_B = 0$ is always an optimal response for Blue, as is $m_R = 0$ for Red. Thus, Blue and Red's utility functions become quadratic functions of only F_B and F_R, where the quadratic form comes from the fact SGVs turn out to be linear in F_B and F_R. This result is formalized in Theorem 5.3, which implies MCS can be ignored in the one-fishery case. It's worth recalling the definition of MCS used in this chapter: measures taken to ensure fishing is occurring where allowed, other than patrols. In other words, $m_B = m_R = 0$ still allows for some degree of strong fisheries management, such as strictly managed quota systems, restrictions on access to fishing gear, and patrols. Not using MCS amounts to giving fishermen with a legal quota a free pass to bring fishermen in through ports, irrespective of where it was caught.

Theorem 5.3

Assume Blue and Red each own one fishery. For a given strategy of the other player, responding with $m_k > 0$ can yield at most equivalent utility as using $m_k = 0$. If costs of MCS are nonzero, then $m_k > 0$ yields strictly less utility than $m_k = 0$.

Proof. See Appendix 5.A.

Theorems 5.1, 5.2, and 5.3 make it relatively straightforward to find an equilibrium for the overall game, allowing for thorough comparative statics. Because the form of the subgame equilibrium depends on which section of the decision space the players' decisions, F_B and F_R, place the game in, the players' utilities likewise depend on this. The key to finding an equilibrium is to find distinct intervals for F_B where Red's optimal response is a non-piecewise function of F_B, and distinct intervals for F_R where Blue's optimal response is a non-piecewise function of F_R, and finally seek pairs of intervals for F_B and F_R where the players' optima intersect. Algorithm 5.2 in Appendix 5.B details how the required intervals are found, and the concept is illustrated in Figure 5.1. In this illustration, an equilibrium exists at $(F_B, F_R) = (878, 1193)$, where the solid grey and black lines intersect.

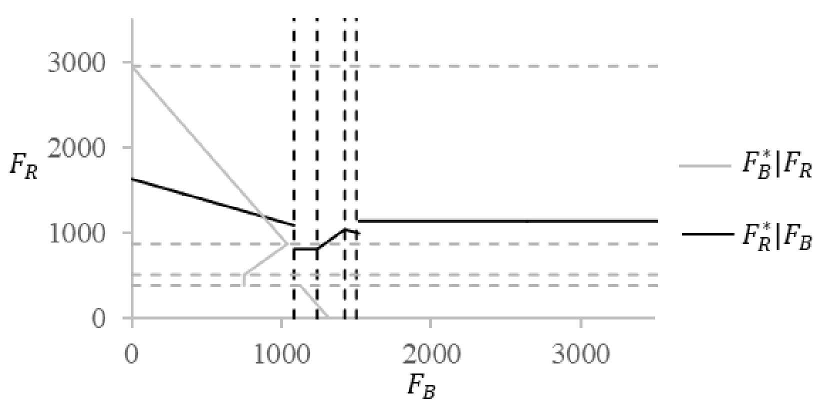

Note. Grey and black dashed lines demarcate intervals where Blue and Red's optimal responses are non-piecewise, respectively.

Figure 5.1. Intervals where optimal responses are non-piecewise.

An important point is that an equilibrium isn't guaranteed to exist for the overall game. Comparative statics are performed in the following example, so in cases where an equilibrium doesn't exist the point that minimizes the value of $\left(u_B^*(F_R) - u_B(F_B, F_R)\right)^2 + \left(u_R^*(F_B) - u_R(F_B, F_R)\right)^2$ is used in its place, where $u_B^*(F_R)$ and $u_R^*(F_B)$ are the optimal utilities Blue and Red can earn when responding to F_R and F_B, respectively. This point is easy to find given the information encapsulated in Algorithm 5.2; one simply iterates through the distinct regions where each $u_B^*(F_R)$ and $u_R^*(F_B)$ are non-piecewise and solves a constrained polynomial optimization. In this chapter the COBYLA method was used (Powell 1998).

5.4.1.1. Example 5.1

Consider the following parameter values, where carrying capacities are measured in millions of metric tons and prices are USD per million metric tons:

$Z_B = Z_R = 2, r_B = r_R = .4, q_B = q_R = .0002, p_B = p_R = \$3{,}000{,}000{,}000$

$c_B = \$200{,}000, P_B = 30, \beta_B = 600$

$c_R = \$120{,}000, P_R = 50, \beta_R = 400.$

The fisheries are symmetric, but Red faces substantially lower opportunity costs and also has more/better patrols and a higher prevalence of maritime militias. These parameters were chosen as it's akin to the ECS, where a less-developed country, China, has lower opportunity costs compared to South Korea and Japan, yet also has made large investments in patrols and maritime militias (Erickson 2018). Table 5.1 shows the equilibrium solution for F_B and F_R and the resulting SGVs, biomasses, and utilities. The solution that would maximize utilities if illegal encroachments were not a possibility ("Legal Optimum") is also provided for comparison.

Table 5.1
Results for Example 5.1

	Equilibrium	**Legal Optimum**
F_B	1042	833
F_R	1269	900
$F_{B,B}$	1042	833
$F_{R,B}$	0	n/a
$F_{B,R}$	103	n/a
$F_{R,R}$	1166	900
x_B	0.8544	1.1667
x_R	0.8344	1.1000
u_B	$325,868,148.15	$416,666,666.67
u_R	$483,023,703.70	$486,000,000.00

Each player is fishing substantially more than the legal optimum. The intuitive explanation for this is that if one player, say Blue, were to use her legal optimum of $F_B = 833$, then Blue biomass is relatively high and Red can issue excessive quotas to induce his fishermen to enter Blue waters. In response, Blue issues excessive quotas to make her waters less desirable to illegal Red fishermen.

Another key takeaway from this example is that Red's utility is largely unaffected by illegal fishing; precisely, it's 99.39% of the legal optimal utility. Blue utility, in contrast, falls to 78.21% of her legal optimum. The disparity comes from the fact Blue fishermen don't fish illegally, because they can't compete with Red fishermen in Red waters due to asymmetries in costs. Overall utility has declined to 89.61% of the legal optimum.

To assess how much each player can gain from cooperation a Nash bargaining problem is solved (Nash 1953). To simplify the analysis, it's assumed each state agrees to implement MCS measures that make illegal fishing unviable, and that this is costless. In Section 5.4.2 the actual costs of MCS will be incorporated and each player will invest in MCS along a continuum. The Nash bargaining problem is:

$$\max_{F_{B,B},\ F_{B,R},\ F_{R,R}} (u_B - u_B^{NC})(u_R - u_R^{NC}) \tag{5.7}$$

$$s.t.\ \ u_B \geq u_B^{NC} \text{ and } u_R \geq u_R^{NC},$$

where u_B^{NC} and u_R^{NC} are the "threat values," defined as the noncooperative equilibrium utilities. The additional simplifying assumption, which will also be dropped in Section 5.4.2, is that $F_{R,B} = 0$. This is sensible, and means that as compensation for implementing MCS, which will disproportionately benefit Blue, Blue must grant Red legal quotas in Blue waters. It's easy to see the optimal value of $F_{R,R}$ is $F_{R,R}^{NBS} = \frac{r_R}{2q_R Z_R}\left(Z_R - \frac{c_R}{pq}\right)$ (as would be the case if Red analyzed his problem in isolation), and that for any value of $F_{B,R}$ the optimal value of $F_{B,B}$ is $F_{B,B}^{NBS} = \frac{r_B}{2q_B Z_B}\left(Z_B - \frac{c_B}{p_B q_B}\right) - \frac{1}{2}F_{B,R}$. Thus, (5.7) can be solved via a univariate search over $F_{B,R}$. The solution is $F_{B,B}^{NBS} = 788$, $F_{B,R}^{NBS} = 90$, and $F_{R,R}^{NBS} = 900$, which yields utilities $u_B^{NBS} = \$372{,}881{,}667$ (a 14.43% gain) and $u_R^{NBS} = \$535{,}770{,}000$ (a 10.92% gain). Notice that 10.25% of all fishing occurring in Blue waters is by Red fishermen. While significant, this figure doesn't seem so high that a perception Blue politicians are giving up their sovereign territory to Red will emerge, and therefore may be politically amenable.

5.4.1.2. Comparative Statics for Example 5.1

Figure 5.2 shows how the results of Example 5.1 are affected by c_R. Notice the same qualitative results hold as c_R spans \$100,000 to \$200,000: Red continues to earn near his legal optimum utility in the noncooperative case, while Blue achieves substantially less; under the Nash bargain, Red achieves significantly more than legality while Blue recoups some, though not all, of her legal utility. This is true even when $c_R = c_B = \$200{,}000$, indicating the influence of Red's superior patrols and maritime militias is not trivial. While not pictured, the bargaining solution is such that $F_{B,R}/\left(F_{B,B} + F_{B,R}\right)$ reaches a peak of 13.33% (when $c_R = \$100{,}000$), which as in Example 5.1 does not seem unreasonably high.

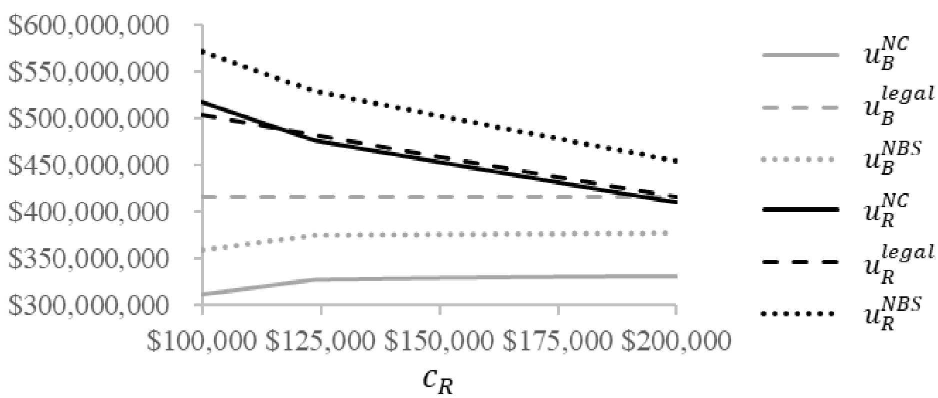

Figure 5.2. Sensitivity of utilities to c_R in Example 5.1.

Table 5.2 provides a comparison of noncooperative and bargaining utilities for additional examples. As when c_R was shocked, Red persistently loses little under noncooperation while Blue loses a lot. In certain cases Red's noncooperative utility is actually significantly higher than his legal optimum; see particular the example of asymmetric carrying capacities, $Z_B = 2$ and $Z_R = 1.5$. Unlike in Figure 5.1, the percentage of Red fishing in Blue waters under the bargain can get very high (as high as 50.63%). This may or may not pose a hindrance in negotiations.

Table 5.2
Additional examples

Parameters	u_B^{NC}	u_B^{NBS}	u_B^{legal}	u_R^{NC}	u_R^{NBS}	u_R^{legal}	$\frac{F_{BR}}{F_{BB}+F_{BR}}$
Baseline (Ex. 5.1)	.33	.37	.42	.48	.54	.49	10.25%
$Z_B = Z_R = 1$	.08	.10	.13	.22	.24	.19	23.79%
$Z_B = Z_R = 3$	.57	.65	.71	.76	.85	.78	8.30%
$p = 2$	.16	.19	.23	.31	.34	.29	16.51%
$p = 2, Z_B = Z_R = 1$	.02	.02	.05	.14	.15	.10	50.63%
$p = 2, Z_B = Z_R = 3$	.33	.37	.42	.48	.54	.49	10.25%
$q = .0001$	.17	.20	.27	.44	.47	.38	23.73%
$q = .0003$	.38	.43	.47	.51	.57	.52	8.25%
$Z_R = 1.5$	.25	.31	.42	.41	.46	.34	25.17%
$Z_R = 1.75$	.29	.34	.42	.45	.50	.41	17.72%

Note. All utilities and prices are measured in billions of USD. Parameters not explicitly stated are the same as in Example 5.1.

5.4.2. Three Fisheries per Player

A limitation of the one-fishery case is it doesn't allow for a comparison of optimal levels of MCS when illegal encroachments are and are not a concern. In an uncongested environment with one fishery, a state clearly has no incentive to spend on MCS since fishermen only have a single fishery to choose from. In reality, states are responsible for multiple fisheries and proper management requires measures to verify fishermen are catching only where allowed. This makes the single-fishery model inadequate to understand the observed poor MCS practices in the ECS. Thus, this subsection considers a scenario where Blue and Red own three fisheries each, shows there is an incentive to use strong MCS if it's assumed illegal encroachments are not a possibility, but that this incentive vanishes in a congested environment where encroachments are a possibility.

With multiple fisheries the problem is not conceptually different, as the same subgame, (5.6), still must be solved. The analytical challenge with multiple fisheries relates to dimensionality, as the proof of Theorem 5.2 involves enumerating all possible combinations of the binary variables, $y_{ij,kl}$ and y_{kl}, and deriving necessary and sufficient conditions and formulae for the subgame equilibrium for each combination. This is burdensome even with two fisheries per player. An alternative solution methodology is to pick random starting values for all F_{kl} and m_k, allow the players to sequentially respond optimally to the other, stopping once neither player can gain more than 0.1% in utility. Blue's optimal responses are found using the subgame (5.6), but with F_{Bl} and m_B now included as decision variables and the objective functions replaced with the utility function (5.4); an analogous statement holds for Red's optimal response. This can lead to a nonconvex program (even when $m_B = 0$), so a global optimization method is required; this chapter used nested partitioning (Shi and Ólafsson 2000).

5.4.2.1. Example 5.2: One Player with Three Fisheries

Analyzing a single-player example (i.e. an uncongested environment) allows one to establish a situation where MCS is a cost-effective measure to deter illegal fishing in the absence of foreign encroachments. The parameters used are similar to Example 5.1, but now there are three fisheries with different prices:

$Z_{Bj} = 2$, $r_{Bj} = .4$, $q_{Bj} = .0002 \ \forall \ j \in \{1,2,3\}$

$p_{B1} = \$3{,}000{,}000{,}000$, $p_{B2} = \$2{,}250{,}000{,}000$, $p_{B3} = \$1{,}500{,}000{,}000$,

$c_B = \$200{,}000$, $P_B = 30$, $\beta_B = 600$

$\beta_m = 6{,}000$, $a_1 = 3{,}393{,}400$, $a_2 = 0.5$.

MCS costs are difficult to estimate as state spending can vary widely, irrespective of MCS quality. Empirical research has nonetheless concluded MCS is cost-effective, so this example uses values of a_1 and a_2 which (a) lead to MCS being used in the one-state case, and (b) lead to spending falling within the range of observed MCS spending among highly-rated countries (see (Mangin et al. 2018) for spending among nations, and (Petrossian 2019) for a summary of U.N. FAO ratings of MCS regimes). This approach is sufficient to illustrate the broader point: MCS is used when a state operates in isolation, but is not used in a congested environment.

Optimizing Blue strategy yields investment in MCS of $MCS_B^{legal} \coloneqq a_1\left(m_B^{legal}\right)^{a_2} = \$50{,}759{,}462$, and fishing quotas of $\left(F_{B1}^{legal}, F_{B2}^{legal}, F_{B3}^{legal}\right) = (876, 776, 632)$. Blue's optimal utility is $u_B^{legal} = \$770{,}036{,}329$. As shown next, in the congested environment neither Blue nor Red will invest in MCS, and Blue's utility will fall substantially.

5.4.2.2. Example 5.3: Two Players, Three Fisheries Each

This example extends Example 5.2, where Red now owns three fisheries which are symmetric to Blue's. Red's costs, patrols, and maritime militias are as they were in Example 5.1: $c_R = \$120{,}000$, $P_R = 50$, $\beta_R = 400$. The iterative search methodology described at the beginning of Section 5.4.2 is used to find the following equilibrium:

$$MCS_B^{NC} = \$0, (F_{B1}^{NC}, F_{B2}^{NC}, F_{B3}^{NC}) = (1220,\ 1002,\ 443)$$

$$MCS_R^{NC} = \$0, (F_{R1}^{NC}, F_{R2}^{NC}, F_{R3}^{NC}) = (1346,\ 1301,\ 1191).$$

Overall, the players are fishing more than would be the case in an uncongested environment; Blue's legal optima was provided in Example 5.2, and Red's is $\left(F_{R1}^{legal}, F_{R2}^{legal}, F_{R3}^{legal}\right) = (964, 869, 733)$. The solution to the subgame at this equilibrium point is shown in Table 5.3, and is unique. Uniqueness was verified by attempting to solve (5.6) multiple times, each time with a single added constraint stating a SGV is above or below its value in Table 5.3; all attempts resulted in an infeasible program. The off-diagonals of Table 5.3 represent illegal fishing and it's seen Red is the only player fishing illegally. Each of these results mirror what was found in the one-fishery case. The new finding is that while MCS was viewed as a valuable investment in the uncongested scenario of Example 5.2, the players now do not invest in MCS.

Table 5.3

Subgame equilibrium for Example 5.3 equilibrium

	Legally Authorized Fishery					
Utilized Fishery	$B1$	$B2$	$B3$	$R1$	$R2$	$R3$
$B1$	1220	0	0	0	128	12
$B2$	0	1002	0	0	0	144
$B3$	0	0	443	0	0	276
$R1$	0	0	0	1346	0	0
$R2$	0	0	0	0	1173	0
$R3$	0	0	0	0	0	759

The utilities realized at this equilibrium are $u_B^{NC} = \$490{,}945{,}571$ and $u_R^{NC} = \$994{,}985{,}413$. Legal utilities are $u_B^{legal} = \$770{,}036{,}329$ and $u_R^{legal} = \$964{,}078{,}848$; the large decline in Blue utility and relative insensitivity of Red utility again mirrors the one-fishery case. To assess the value of cooperation, the Nash bargaining problem (5.7) is again solved, but now MCS costs are accounted for and each player is allowed to fish legally in each fishery of the other. Nested partitioning was used to arrive at the solution $MCS_B^{NBS} = \$54{,}733{,}849$, $MCS_R^{NBS} = \$56{,}347{,}437$, and the fishing levels detailed in Table 5.4. The corresponding utilities are $u_B^{NBS} = \$618{,}337{,}318$ (a gain of 25.95%) and $u_R^{NBS} = \$1{,}160{,}859{,}577$ (a gain of 16.67%).

Table 5.4

Nash bargaining quotas for Example 5.3

	Fishery					
Player	$B1$	$B2$	$B3$	$R1$	$R2$	$R3$
Blue	800	726	91	14	59	10
Red	58	86	640	896	836	720

As in the one fishery case, it's useful to examine how many quotas are being granted for foreign fishermen to fish in domestic waters. In all but fishery $B3$, the percentage of foreign fishing is small. Fishery $B3$ has in effect been given to Red as payback for cooperating (87.55% of all fishing here is by Red fishermen). This may cause political problems, but examining this is beyond the scope of this chapter.

5.4.2.3. Additional Examples with Three Fisheries per Player

The results of Example 5.3 are generalized in Table 5.5, which compares noncooperative equilibrium utilities, Nash bargaining utilities, and legal utilities for alternative values of P_B and c_R. Note that because P_B only enters the model in the term $\beta_R P_B$, shocking P_B can also be used to assess the sensitivity to β_R. Nash bargaining MCS expenditures are also listed. All examples exhibit the same pattern as past findings. Of note, Blue's utility increases substantially under both noncooperation and bargaining as P_B increases from the baseline value of 30 patrols craft per 100,000 square kilometers, to 50. The gain from increasing P_B from 50 to 100 is also significant. Of course, viewing this decision in isolation neglects the possibility Red may respond by investing in more maritime militias, offsetting the increase in Blue patrols. The strategic interaction between investments in patrols and maritime militias is left as a point of future research.

Table 5.5
Utilities for additional examples with multiple fisheries

Parameters	u_B^{NC}	u_B^{NBS}	u_B^{legal}	u_R^{NC}	u_R^{NBS}	u_R^{legal}	MCS_B^{NBS}	MCS_R^{NBS}
	billions of USD						*millions of USD*	
Base (Ex. 5.3)	.49	.62	.7700	.99	1.16	.9641	54.73	56.35
$P_B = 50$	.59	.67	.7715	.97	1.07	.9641	46.53	48.25
$P_B = 100$	.63	.71	.7748	.91	1.02	.9641	34.12	60.63
$c_R = \$150{,}000$	.58	.69	.7700	.86	.98	.8887	44.08	52.17
$c_R = \$200{,}000$	.71	.76	.7700	.72	.77	.7704	50.77	51.36

5.5. Conclusion and Future Work

This chapter modeled a scenario akin to the ECS, where two maritime nations in close proximity faced the decisions of how many legal fishing quotas to issue and how much to invest in MCS. Consistent with observed behavior in the ECS, the model predicts illegal encroachments on account of excessive issuance of quotas. Also consistent with observed behavior, each state underinvests in MCS. The chapter thus provides a rational explanation for illegal encroachments and substandard MCS beyond the typical explanations of uncontrollable criminal activity and a lack of political will to implement strong MCS. A bargaining problem was also solved which quantified the substantial gains to be had from cooperation. In the bargain, the player with lower costs, either by virtue of lower opportunity costs or more patrols and maritime militias, is seen to achieve more than the optimal utility he could achieve in the absence of a nearby player. This gain is realized through legal quotas for the low-cost player in the high-cost player's waters.

The Gordon-Schaeffer fisheries model with logistics growth and a constant-effort harvesting rates was used throughout this chapter. Future research ought to perform similar analysis with alternative fisheries models. Additional complicating factors should also to be assessed, such as migration of species, discounted payoffs, and stochastic growth and harvest rates. Problems with many more than three fisheries, and more than two players, are perhaps the most critical point of future research as such a scenario is required to truly model the ECS, and thus use the model for policy recommendations. All of these model extensions will require sophisticated computational techniques.

A few other lines of future work are pertinent. As mentioned in Section 5.2, an empirical analysis of the effectiveness of maritime militias is lacking in the literature. Such a study would improve the analysis presented here by giving evidence-based assessments of β_B and β_R. Accounting for illegal third-parties, such as long-distance fishermen entering the ECS and unloading their catch at far-off ports would also add realism to the model. Lastly, congested maritime environments where there is not a legal delineation of fishing rights is an important area of study. Such is the situation in the SCS, for example, where the consequences of opposing nations' patrol craft confronting one another becomes an important consideration that wasn't relevant in this chapter. Not only do interacting patrols likely reduce the deterrent effect of patrols (as fishermen can sound a distress call if confronted by opposition patrols), but the threat of escalation to greater conflict may increase.

Appendix 5.A. Proofs of Theorems 5.1-3

This appendix restates and proves Theorems 5.1, 5.2, and 5.3. Throughout, notation is used that assumed each player owns only one fishery. To support the proofs, a useful, nonrestrictive assumption is stated.

Assumption 5.1. Based on the statistical improbability of the contrary, it's assumed open access levels (the levels of biomasses at which rents are 0) are never exactly identical for Blue and Red fishermen. That is, $\frac{c_B}{p_B q_B} \neq \frac{c_R + \beta_R P B_B}{p_B q_B}$ and $\frac{c_B + \beta_B P B_R}{p_R q_R} \neq \frac{c_R}{p_R q_R}$.

5.A.1. Proof of Theorem 5.1

Theorem 5.1. Existence of a unique subgame equilibrium.

When Blue and Red each own only one fishery, any instantiation of the game's parameters and choice of the overall game's decision variables yields a unique subgame equilibrium (SGE).

Proof.

It's first be proved that for any instantiation, a SGE exists. Consider the following algorithm which takes the parameters and (F_B, F_R) as given and seeks values of the SGVs:

Algorithm 5.1. Finding a SGE.

1. Initiate $F_{BB} = F_{RB} = F_{BR} = F_{RR} = 0$.

2. Determine the set of SGVs, F', to increase in the next iteration of the algorithm, per below.

Define the rents produced by each SGV as $\pi_{BB} = p_B q_B x_B - c_B$, $\pi_{RB} = p_B q_B x_R - c_B - \beta_B PB_R$, $\pi_{BR} = p_R q_R x_B - c_R - \beta_R PB_B$, and $\pi_{RR} = p_R q_R x_R - c_R$; e.g. the SGV F_{BR} produces rent equivalent to R_{BR}. Define the maximum achievable rent as:

$$\pi' = \max \begin{cases} \pi_{BB}, \pi_{RB}, 0;\ \text{if } F_{BR} + F_{RR} = F_R, F_{BB} + F_{RB} < F_B \\ \pi_{BR}, \pi_{RR}, 0;\ \text{if } F_{BR} + F_{RR} < F_R, F_{BB} + F_{RB} = F_B \\ 0;\ \text{if } F_{BR} + F_{RR} = F_R, F_{BB} + F_{RB} = F_B \\ \pi_{BB}, \pi_{RB}, \pi_{BR}, \pi_{RR}, 0;\ \text{otherwise} \end{cases}. \tag{5.8}$$

If $\pi' = 0$, stop. Otherwise, include all SGVs in F' whose corresponding rent equals π', and can be increased without decreasing another SGV; that is, only F_{BB} and F_{RB} may only be in F' if $F_{BB} + F_{RB} < F_B$ and $F_{BR} + F_{RR} = F_R$, and similarly for F_{BR} and F_{RR}.

Before describing steps 3 through 6, note the following properties are maintained throughout the algorithm:

Algorithm 5.1, Property 1. SGVs are never allowed to go below 0, nor are the total quotas used by Blue (Red) allowed to exceed F_B (F_R), so constraints (5.6b) and (5.6e) are satisfied throughout the algorithm.

Algorithm 5.1, Property 2. If a SGV is greater than 0, it must produce positive rent which is at least as great as the corresponding player's rent in the other waters. Therefore, constraint (5.6c) governing equality of rent are satisfied throughout the algorithm.

3. If $|F'| = 1$, then WLOG assume $F' = \{F_{BB}\}$. Note this implies $F_{RB} = 0$; it's known $R_{RB} < R_{BB} = R'$ (otherwise F_{RB} would be in F'), and because the equality of rent constraints are maintained throughout the algorithm, $F_{RB} = 0$.

3.1. If $F_{BR} = 0$, increase F_{BB} until $\pi_{BB} = 0$, $F_{BB} = F_B$, or π_{BB} equals one of π_{RB}, π_{BR}, or π_{RR} (whichever occurs first). Return to step 2.

3.2. If $F_{BR} > 0$ and $\pi_{BR} > \pi_{RR}$, increase F_{BB} until $\pi_{BB} = 0$, $F_{BB} = F_B$, $\pi_{BR} = 0$, $\pi_{BR} = \pi_{RR}$, or π_{BB} equals one of π_{RB} or π_{BR} (whichever occurs first). Return to step 2.

3.3. If $F_{BR} > 0$ and $\pi_{BR} = \pi_{RR} > 0$, increase F_{BB} at rate 1, decrease F_{BR} at rate $\Delta = \frac{Z_B q_B r_B r_R}{r_B(Z_B q_R r_R + Z_R q_R r_B)}$, and increase F_{RR} at rate Δ until $\pi_{BB} = 0$, $F_{BB} = F_B$, $F_{BR} = 0$, $\pi_{BR} = \pi_{RR} = 0$, or $\pi_{BB} = \pi_{BR} = \pi_{RR}$ (whichever occurs first). Return to step 2. Note: Δ is defined so equality of Red rents is preserved; also note the total quotas being used by Red fishermen does not change.

3.4. If $F_{BR} > 0$ and $\pi_{BR} = 0$, increase F_{BB} at rate 1 and decrease F_{BR} at rate $\Delta = \frac{q_B}{q_R}$ until $F_{BB} = F_B$ or $F_{BR} = 0$ (whichever occurs first). Return to step 2. Δ is defined so Blue biomass is not changed; Blue fishermen merely replace Red ones.

4. If $|F'| = 2$ and both SGVs in F' represent the same player, then assume WLOG $F' = \{F_{BB}, F_{RB}\}$. Note: $\pi_{RR} > \pi_{BR}$, since Blue is achieving equality of rent; hence, $F_{BR} = 0$.

4.1. If $F_{RR} = 0$, increase F_{BB} at rate 1 and F_{RB} at rate $\Delta = \frac{r_R z_B}{r_B z_R}$, until $\pi_{BB} = \pi_{RB} = 0$, $F_{BB} + F_{RB} = F_B$, or $\pi_{BB} = \pi_{RB} = \pi_{RR}$ (whichever occurs first). Return to step 2. Δ was derived so equality of Blue rents is maintained.

4.2. If $F_{RR} > 0$ and $\pi_{RR} > 0$, increase F_{BB} at rate 1 and F_{RB} at rate $\Delta = \frac{r_R z_B}{r_B z_R}$ until $\pi_{BB} = \pi_{RB} = 0$, $F_{BB} + F_{RB} = F_B$, $\pi_{RR} = 0$, or $\pi_{BB} = \pi_{RB} = \pi_{RR}$ (whichever occurs first). Return to step 2. Δ was again defined to preserve equality of Blue rents.

4.3. If $F_{RR} > 0$ and $\pi_{RR} = 0$, increase F_{RB} at rate 1 and decrease F_{RR} at rate $\Delta = \frac{q_B}{q_R}$ until $F_{BB} + F_{RB} = F_B$ or $F_{RR} = 0$. Return to step 2. Δ is defined so Red biomass is not changed; Blue fishermen merely replace Red ones.

5. If $|F'| = 2$ and the SGVs in F' represent different players, there are two general cases to consider.

5.1. Assume each player is fishing domestically, so that $F' = \{F_{BB}, F_{RR}\}$; note this implies $F_{RB} = F_{BR} = 0$, so the equality of rent constraints aren't violated. Increase F_{BB} at rate 1 and F_{RR} at rate $\Delta = \frac{r_R p_B q_B^2 Z_B}{r_B p_R q_R^2 Z_R}$ until $\pi_{BB} = \pi_{RR} = 0$, $F_{BB} = F_B$, $F_{RR} = F_R$, or $\pi_{BB} = \pi_{RR}$ equals one of π_{RB} or π_{BR} (whichever occurs first). Return to step 2. Δ was determined so $\pi_{BB} = \pi_{RR}$ is maintained.

5.2. Assume both players are fishing in one fishery. WLOG, let $F' = \{F_{BB}, F_{BR}\}$. Unless $p_B q_B = p_R q_R$ this represents an infinitesimally small portion of x_B-space: $p_B q_B x_B - c_B = p_R q_R x_B - c_R - \beta_R PB_B \rightarrow x_B = \frac{c_B - c_R - \beta_R PB_B}{p_B q_B - p_R q_R}$. As soon as x_B decreases by an infinitesimal amount, Blue fishermen are receiving higher rent in Blue waters if $p_B q_B < p_R q_R$, and Red is receiving higher rent otherwise.

5.2.1. If $p_B q_B = p_R q_R$, then as x_B changes Blue and Red rents change by the same amount. Increase F_{BB} and F_{BR} at the same rate until $\pi_{BB} = \pi_{BR} = 0$, $F_{BB} = F_B$, $F_{BR} = F_R$, or $\pi_{BB} = \pi_{BR} = \pi_{RR}$ (whichever occurs first). Return to step 2.

5.2.2. If $p_B q_B \neq p_R q_R$, WLOG assume $p_B q_B > p_R q_R$. Increase F_{BR} until $\pi_{BB} = 0$, $F_{BR} = F_R$, or $\pi_{BR} = \pi_{RR}$ (whichever occurs first). Return to step 2.

6. If $|F'| = 3$, assume WLOG $F' = \{F_{BB}, F_{RB}, F_{RR}\}$. Note it's impossible to have the player who's only fishing in one fishery be fishing illegally. For example, if Blue's rents are equal, then Red's is greater in Red waters and hence F_{BR} cannot be in F': $p_B q_B x_B = p_B q_B x_R - \beta_B PB_R \rightarrow p_R q_R x_B - \beta_R PB_B < p_R q_R x_R$. It's also known $F_{BR} = 0$. Similarly to step 5.2 this represents an infinitesimally small portion of (x_B, x_R)-space: $F_{BB}, F_{BR} \in F' \rightarrow p_B q_B x_B - c_B = p_R q_R x_B - c_R - \beta_R PB_B \rightarrow x_B = \frac{c_B - c_R - \beta_R PB_B}{p_B q_B - p_R q_R}$; and $F_{BB}, F_{RB} \in F' \rightarrow p_B q_B x_B = p_B q_B x_R - \beta_B PB_R \rightarrow x_R = x_B + \frac{\beta_B PB_R}{p_B q_B} = \frac{c_B - c_R - \beta_R PB_B}{p_B q_B - p_R q_R} + \frac{\beta_B PB_R}{p_B q_B}$.

6.1. If $p_B q_B = p_R q_R$, then π_{RB} and π_{RR} decline at the same rate as x_R declines. Increase F_{RB} and F_{RR} at rate 1, and increase F_{BB} at rate $\Delta = \frac{r_B Z_R (q_B + q_R)}{r_R Z_B q_B}$, until $F_{RR} = F_R$, $F_{BB} + F_{RB} = F_B$, or $\pi_{BB} = \pi_{RB} = \pi_{RR} = 0$ (whichever occurs first). Return to step 2. Δ was derived to preserve equality of rent.

6.2. If $p_B q_B > p_R q_R$, then π_{RB} declines faster than π_{RR} as x_R declines. Increase F_{RR} at rate 1 and increase F_{BB} at rate $\Delta = \frac{r_B Z_R q_R}{r_R Z_B q_B}$, until $F_{RR} = F_R$, $F_{BB} + F_{RB} = F_B$, or $\pi_{BB} = \pi_{RB} = 0$ (whichever occurs first). Return to step 2. Δ preserves equality of rent for Blue.

6.3. If $p_B q_B < p_R q_R$, then π_{RR} declines faster than π_{RB} as x_R declines. Increase F_{RB} at rate 1 and increase F_{BB} at rate $\Delta = \frac{r_B z_R}{r_R z_B}$, until $F_{BB} + F_{RB} = F_B$ or $\pi_{RR} = 0$ (whichever occurs first). Return to step 2. Δ was derived to preserve equality of rent for Blue.

The proof this algorithm always produces a SGE is clear after noting two additional properties to accompany Properties 1 and 2. Rereading the algorithm, the following properties are apparent:

Algorithm 5.1, Property 3. π' is nonincreasing for every step. More precisely, it's either constant (steps 3.4 and 4.3) or decreasing linearly (all other steps). Steps 3.4 and 4.3 eventually terminate and result in either: $\pi' = 0$ on account of all profitable quotas being used; or performing step 3.1 or 4.1 in the next iteration. Therefore π' will eventually reach 0.

Algorithm 5.1, Property 4. $\pi' = 0$ implies either all quotas are exhausted, or neither player can achieve positive rent with their remaining quotas, and hence constraint (5.6d) is satisfied at the algorithm's termination.

Properties 1 through 4 combined mean all subgame constraints are satisfied at the algorithm's termination, and hence it produces a SGE.

To complete the proof of Theorem 5.1, it's now shown whenever a SGE exists, it must be unique. This is shown in two parts: first, if there exists a SGE resulting in biomasses x_B^{SGE} and x_R^{SGE}, it's shown any SGE must yield these biomasses; next, it's shown biomasses at SGE are sufficient to uniquely determine the values of the SGVs.

Assume Algorithm 5.1 was used to produce a SGE with biomasses x_B^{SGE} and x_R^{SGE}. Now consider an alternative SGE; its biomasses and SGVs will be referenced using the "prime" superscript ($'$) to differentiate between the original SGE, superscripted by SGE. Assume one of the biomasses has increased; WLOG, $x_B' > x_B^{SGE}$. This means at least one player, WLOG Blue, was previously fishing in Blue waters and is now fishing less in Blue waters: $0 \leq F_{B,B}' < F_{B,B}^{SGE}$. If Blue fishermen are not fishing in Red waters under the new SGE ($F_{R,B}' = 0$), then they're using all available quotas in Blue waters: $F_{B,B}' = F_B$; otherwise constraints (5.6c) in the subgame would be violated stating Blue doesn't leave profitable quotas unused. This contradicts $F_{B,B}' < F_{B,B}^{SGE} \leq F_B$, so assume instead Blue fishermen are fishing in Red waters. $F_{R,B}' > 0$ implies x_R has also increased, since by (5.6c) $p_B q_B x_B^{SGE} \geq p_R q_R x_R^{SGE} - \beta_B P_R \rightarrow p_B q_B x_B' > p_R q_R x_R^{SGE} - \beta_B P_R$ and $p_B q_B x_B' \leq p_R q_R x_R' - \beta_B P_R \rightarrow x_R^{SGE} < x_R'$. Since both biomasses have increased, at least one player's fishermen are using less overall quotas, hence they have quotas to use which are profitable at these higher levels of biomass (otherwise, the original SGE would violate (5.6c)), thus (5.6d) is violated and this is not a SGE. In sum, there can be no SGE with $x_B > x_B^{SGE}$, and by analogy none with $x_R > x_R^{SGE}$.

Assume instead one of the biomasses has declined; WLOG, $x_B' < x_B^{SGE}$. At least one of the players, WLOG Blue, must be using more quotas in Blue waters: $F_{B,B}' > F_{B,B}^{SGE} \geq 0$. She would have used this larger amount of quotas in Blue waters under the original SGE so as not to violate (5.6e), *unless* $F_{R,B}^{SGE} > 0$ and Blue fishermen were earning greater or equal rent compared to Blue waters. This implies $x_R' < x_R^{SGE}$, as otherwise Blue would now be receiving strictly greater rent in Red waters, entailing $F_{B,B}' = 0$. Because both biomasses have declined, at least one player's fishermen are using more overall quotas, implying in the original SGE they had additional quotas to use which were profitable at higher biomasses (if not, (5.6c) would be violated in the new SGE), hence (5.6d) was violated by the original SGE, a contradiction. In sum, under any SGE, Blue biomass must be x_B^{SGE}, and by analogy Red must be x_R^{SGE}.

Knowing any SGE must have the same levels of biomass, the proof SGEs are unique is concluded by showing biomasses can uniquely determine the values of the SGVs. First note biomasses can immediately identify which SGVs must be 0. For example, if Red biomass is less than $\frac{c_B+\beta_B p_R}{p_R q_R}$ (Blue's so-called open access level in Red waters), then $F_{R,B} = 0$; similar statements hold for $F_{B,B}$, $F_{B,R}$, and $F_{R,R}$. If $x_B^{SGE} > x_R^{SGE} - \frac{\beta_B p_R}{p_R q_R}$, then Blue fishermen earn greater rent in Blue water and $F_{R,B} = 0$; similar statements hold for $F_{B,B}$, $F_{B,R}$, and $F_{R,R}$. Now consider three possibilities: only one SGV is potentially nonzero; only two are potentially nonzero; and three are potentially nonzero. Note four nonzero SGVs is impossible, as either $x_B^{SGE} > x_R^{SGE} - \frac{\beta_B p_R}{p_R q_R}$, implying

$F_{R,B} = 0$, or $x_B^{SGE} < x_R^{SGE} - \frac{\beta_B P_R}{p_R q_R}$, implying $F_{B,B} = 0$, or $x_B^{SGE} = x_R^{SGE} - \frac{\beta_B P_R}{p_R q_R}$, implying $x_B^{SGE} - \frac{\beta_R P_B}{p_B q_B} < x_R^{SGE}$ and $F_{B,R} = 0$.

If only one SGV may be nonzero, assume WLOG this is $F_{B,B}$. It's apparent this has a unique value; simply increase $F_{B,B}$ until $x_B = x_B^{SGE}$. If two nonzero SGVs are possible there are three general cases. First, assume one player may be fishing in one fishery, and the other player in the other; WLOG, $F_{B,B}$ and $F_{R,R}$ may be nonzero. As when only one SGV could be nonzero, the unique SGE is clear: increase $F_{B,B}$ until $x_B = x_B^{SGE}$, and do likewise for $F_{R,R}$. If instead one player is fishing in both waters, assume WLOG $F_{B,B}$ and $F_{R,B}$ may be nonzero. Again, the unique solution mandates $F_{B,B}$ and $F_{R,B}$ are increased until stocks are depleted to x_B^{SGE} and x_R^{SGE}. Lastly, assume both players may be fishing in one fishery; WLOG, $F_{B,B}$ and $F_{B,R}$ may be nonzero. It's known $x_B^{SGE} \geq \frac{c_B}{p_B q_B}$ and $x_B^{SGE} \geq \frac{c_R + \beta_R P B_B}{p_B q_B}$, and by Assumption 1 one of these inequalities must be strict; WLOG, assume $x_B^{SGE} > \frac{c_R + \beta_R P B_B}{p_B q_B}$. By subgame constraint (5.6d), Red fishermen must use all quotas, $F_{B,R} = F_R$. To uniquely determine $F_{B,B}$, simply increase it until $x_B = x_B^{SGE}$. The final case assumes three SGVs may be nonzero; WLOG, $F_{B,B}$, $F_{B,B}$, and $F_{R,R}$. $F_{R,R}$ must be that which depletes Red biomass to x_R^{SGE}. From here, the same logic can be applied as when only $F_{B,B}$ and $F_{B,R}$ were allowed to be nonzero to determine their unique values; the one amendment is that $F_{B,R} = F_R - F_{R,R}$, rather than

$F_{B,R} = F_R$. In sum, knowing the biomasses at the subgame equilibrium allows the SGVs to be uniquely determined.

To summarize the proof of Theorem 5.1: (i) by Algorithm 5.1 and the associated proof, a SGE always exists; (ii) if there exists a SGE resulting in biomasses x_B^{SGE} and x_R^{SGE}, then any SGE must result in these biomasses; (iii) SGE biomasses are sufficient to uniquely determine the SGVs, and hence a unique SGE always exists.

5.A.2. Proof of Theorem 5.2

Theorem 5.2. Formulae for the unique subgame equilibrium via a partitioning of the parameter and decision space.

Assume Blue and Red each own one fishery. The parameter and decision variable space can be partitioned, such that in each region of the partition analytical formulae exist for the subgame variables.

Proof.

The proof uses Theorem 5.1, as well as the notion that subgame equilibria can be categorized into various "forms." The form of a subgame equilibrium defines which SGVs are nonzero and whether each player's quotas are exhausted. For example, one form of a SGE could be $F_{B,B}, F_{B,R}, F_{R,R} > 0$, $F_{R,B} = 0$, $F_{B,B} = F_B$, and $F_{B,R} + F_{R,R} = F_R$. For each possible form, it's straightforward to derive necessary and sufficient conditions for the SGE to have that form, as is deriving the analytical formulae for the SGVs. Continuing the example form just mentioned, the formulae for the SGVs are derived by noting that since $F_{B,R}, F_{R,R} > 0$, Red's rents must equate (by (5.6c)): $p_B q_B Z_B \left(1 - \frac{q_B}{r_B}\left(F_B + F_{B,R}\right)\right) - \beta_R P_B - \beta_m m_R = p_R q_R Z_R \left(1 - \frac{q_R}{r_R} F_{R,R}\right)$. Combined with $F_{B,R} + F_{R,R} = F_R$, this yields a linear system with two equations and two unknowns, which has solution:

$$F_{B,R} = \frac{-r_R p_B q_B^2 Z_B F_B + r_B p_R q_R^2 Z_R F_R + r_R r_B [p_B q_B Z_B - p_R q_R Z_R - \beta_R P_B - \beta_m m_R]}{r_B p_R q_R^2 Z_R + r_R p_B q_B^2 Z_B} \qquad (5.9)$$

$$F_{R,R} = \frac{r_R p_B q_B^2 Z_B F_B + r_R p_B q_B^2 Z_B F_R - r_R r_B [p_B q_B Z_B - p_R q_R Z_R - \beta_R P_B - \beta_m m_R]}{r_R p_B q_B^2 Z_B + r_B p_R q_R^2 Z_R}. \qquad (5.10)$$

$F_{B,B} = F_B$ and $F_{R,B} = 0$ is given. Identifying necessary and sufficient conditions is a simple matter of checking which of constraints (5.6c) and (5.6d) are applicable, and imposing strict positivity constraints. In this example, those conditions are:

Positivity of $F_{B,R}$: $0 < F_{B,R}$ (5.11)

Positivity of $F_{R,R}$: $0 < F_{R,R}$ (5.12)

Blue fishermen achieve nonnegative rent in Blue waters: $\pi_{B,B} \geq 0$ (5.13)

Red fishermen's rents are nonnegative: $\pi_{B,R} \geq 0$. (5.14)

By virtue of (5.9) and (5.10), constraints (5.11) – (5.14) are linear in F_B and F_R. Also note (5.13) ensures $\pi_{B,B} \geq \pi_{R,B}$, since by construction $\pi_{B,R} = \pi_{R,R} \rightarrow p_B q_B x_B - \beta_R P_B - \beta_m m_R = p_R q_R x_R \rightarrow p_B q_B x_B > p_R q_R x_R - \beta_B P_R - \beta_m m_B \rightarrow \pi_{B,B} > \pi_{R,B}$. (5.14) ensures $\pi_{R,R} \geq 0$, since $\pi_{B,R} = \pi_{R,R}$. This same procedure for finding the subgame equilibrium formula and necessary and sufficient conditions can be done for all possibly forms (there are a total of 29), and is omitted for space. With this notion in hand the proof of the theorem is as follows:

- For a subgame equilibrium of a particular form: (i) the SGVs must solve an associated system of linear equations and unknowns with a unique solution; (ii) necessary and sufficient conditions exist for the SGE to have that form.
- Given values of the parameters and decision variables, produce a SGE via Algorithm 5.1. By Theorem 5.1, this is the unique SGE.
- The SGE will have some form. By the necessity of the conditions for each form, those conditions must hold, and hence the 29 sets of conditions encompass the entire parameter/decision space.
- If any other set of conditions held, then by sufficiency a SGE of another form would hold, which is impossible because SGEs are unique.
- Therefore, for any instantiation of the parameters and decision variables the conditions of one and only one form may hold. This completes the proof.

5.A.3. Proof of Theorem 5.3

Theorem 5.3

Assume Blue and Red each own one fishery. For a given strategy of the other player, responding with $m_k > 0$ can yield at most equivalent utility as using $m_k = 0$. If costs of MCS are nonzero, then $m_k > 0$ yields strictly less utility than $m_k = 0$.

Proof.

WLOG assume Blue is using non-zero MCS: $m_B > 0$. Unilaterally changing strategy to $m_B = 0$ only affects the costs imposed on Blue fishermen in Blue waters. If $F_{R,B}$ does not change at subgame equilibrium on account of the change in m_B (because it was not previously profitable to illegally fish in Red waters, and is still not), then all SGVs are unchanged and Blue's utility is either the same (if MCS is costless) or has increased (if MCS costs money). If $F_{R,B}$ does change, then it increases due to the reduction in costs. This means either: (i) Blue extracts rent from Red waters, improving her utility; or (ii) Blue fishermen have diverted from Blue to Red waters. In the latter case, Blue can simply issue more quotas to replace those who diverted to Red waters, establishing a subgame equilibrium equivalent to that when $m_B > 0$, but with a higher value of $F_{R,B}$. By the previous logic, Blue utility has increased. This completes the proof.

Appendix 5.B. Algorithm to Find Intervals where Optimal Responses are Non-Piecewise

This appendix provides an annotated algorithm to identify intervals in F_B-space where Red's optimal response is a non-piecewise function, and likewise for Blue's optimal response in intervals of F_R-space. The algorithm works by finding which conditions partitioning the subgame equilibrium forms (see Appendix 5.A.2) could be satisfied, given the model parameters, and then identifies key intersections of the conditions.

Algorithm 5.2. Identifying the critical intervals of F_B and F_R.

Given values of Z_B, Z_R, c_B, c_R, β_B, β_R, PB_B, PB_R, r_B, r_R, q_B, q_R, p_B, and p_R:

Eliminate infeasible SGE forms (as defined in Appendix 5.a) and unnecessary conditions, based on parameter values

1. Remove impossible forms. Particularly, identify all forms with conditions $a > b$ where the parameters lead to $a < b$, and rule these forms out.

2. Remove unnecessary conditions. Particularly, for each remaining form identify the greatest lower bound on each F_B, F_R, $q_B F_B + q_R F_R$, and $r_R Z_B q_B F_B - r_B Z_R q_R F_R$ (including the lower bound of 0) and eliminate all but the greatest lower bound. Do the same for the smallest upper bound, including infinity. Rule out any form where the greatest lower bound exceeds the smallest upper bound.

3. Treat all remaining conditions as equalities and label them $c_1, c_2, \ldots, c_N$, where N is the number of unique conditions after converting them to equalities.

Identify where the remaining conditions intersect

4. Loop through conditions c_1 through c_{N-1}.

    4.1. Call the current condition c_i. Loop through conditions c_{i+1} to c_N.

        4.1.1. Call the second condition c_j. If conditions c_i and c_j intersect, record the intersection (both the values of F_B and F_R, and which conditions cross at this point). Note: in this chapter there is at most one intersection per pair of conditions, since all conditions are linear.

Order the conditions by F_R-value as a function of F_B, and vice versa

5. For all conditions that are not parallel to the F_R-axis (i.e. $F_B = a$, a constant), determine the order of these conditions, from smallest to largest, at $F_B = 0$. Store the order associated with $F_B = 0$.

6. Loop through the values of F_B at which conditions intersect, from smallest to largest, found in step 4.

   6.1. Call the current F_B value F_B', and associate the previous ordering of the conditions with F_B'. This ordering will be modified in step 6.2.

   6.2. Loop through all F_R values such that (F_B', F_R) is an intersection point for conditions, starting with the smallest such value.

      6.2.1. For each value of F_R, reorder the conditions intersecting at (F_B', F_R) from smallest to largest to the immediate right of F_B'. Note: because the conditions are linear, this can be done by ordering the conditions according to their slopes.

7. Steps 5 and 6 have created an ordering, from smallest to largest, of the conditions' F_R-values when F_B is within a specified interval. Repeat steps 5 and 6 with F_R and F_B inverted to create an ordering by F_B-value, given F_R.

Create a sequence of critical conditions and SGE forms between values of F_B and F_R where intersections occur

8. Again loop through the values of F_B at which conditions intersect, from smallest to largest, found in step 4.

    8.1. Call the current F_B value F_B'. Loop through the conditions c_1 to c_N to the immediate right of F_B', from smallest to largest, as found in steps 5 through 6.

        8.1.1. Call the current condition c_i. Record the SGE form directly below condition c_i, and label it r_i.

        8.1.2. If $r_i \neq r_{i-1}$, then add condition c_{i-1} to a list of critical conditions associated with F_B', at which the SGE form changes. Also add form r_{i-1} to a similar list.

9. Step 8 provides a list of conditions, for any interval of F_B-space, at which Red's decision F_R causes the game to move from one SGE form to another. Repeat step 8 with F_R and F_B inverted to create a similar list for conditions on F_B, given F_R.

Identify values of F_B where the sequences of critical conditions found in step 8 changes, indicating the possibility *the formula for the optimal response will change*

10. Begin a list of left-endpoints for values of F_B, demarcating intervals. Initiate the first element as $F_B = 0$. Loop through the values of F_B identified in step 4 (from smallest to largest).

10.1. Call the current F_B value F_B^i. If the sequence of critical conditions associated with F_B^i differs from that associated with F_B^{i+1}, as identified in step 8, add F_B^{i+1} to the list of left-endpoints.

Now, for each interval in F_B-space, find constrained optimal responses where Red's response is constrained to be between consecutive critical conditions. Add additional values of F_B to the list identified in step 10 in order to eliminate any piecewise response function.

11. Loop through the values of F_B identified in step 10.

11.1. Call the current F_B value F_B^i. Loop through critical conditions associated with F_B^i.

11.1.1. Call the current critical condition c_j, and note this represents the upper bound on F_R before moving the game into a new SGE form, while c_{j-1} is the lower bound. Call the form the game is in while $c_{j-1} \leq F_R \leq c_j$ and $F_B^i < F_B < F_B^{i+1}$, v_{ij}. Call the formula for Red's optimal response between these bounds f_{ij}. Determine whether f_{ij} intersects either c_j or c_{j-1} at any point between F_B^i and F_B^{i+1}.

11.1.1.1. If there is no intersection, perform the following:

11.1.1.1.1. If $c_{j-1} < f_{ij} < c_j$, then f_{ij} is the optimal response that keeps the game in form v_j. Record f_{ij}

as the optimal response between F_B^i and F_B^{i+1}, constrained to $c_{j-1} \leq f_{ij} \leq c_j$.

11.1.1.1.2. If $f_{ij} < c_{j-1}$, then the optimal response that keeps the game in form v_{ij} is c_{j-1}. This is so because Red's utility function is concave. Record c_{j-1} as the optimal response between F_B^i and F_B^{i+1}, constrained to $c_{j-1} \leq f_{ij} \leq c_j$.

11.1.1.1.3. If $c_j < f_{ij}$, then the optimal response that keeps the game in form v_{ij} is c_j. Again, this is so because Red's utility function is concave. Record c_j as the optimal response between F_B^i and F_B^{i+1}, constrained to $c_{j-1} \leq f_{ij} \leq c_j$.

11.1.1.2. If there is an intersection, identify the smallest, insert that value of F_B into the list identified in step 10 immediately following F_B^i, and go to step 11.1.1.1.1. Note the series of critical conditions associated with the inserted value is identical to that associated with F_B, and f_{ij} does not intersect c_{j-1} or c_j between F_B^i and the new value of F_B^{i+1} (because F_B^{i+1} is by construction the smallest such point).

Now determine the unconstrained optimal response for each interval in F_B-space, adding additional values of F_B to the list created in steps 10 and 11 to ensure no piecewise optima exist

12. Loop through the list of left-endpoints identified in steps 10 and 11.

    12.1. Call the current F_B value F_B^i, call its associated critical conditions c_{i1}, c_{i2}, …, c_{iN}, and call the utility function computed at the optimal response between critical conditions c_{j-1} and c_j u_{ij}^*. Let the incumbent optimal utility be $u_{iI}^* = u_{i1}^*$. Loop through c_{i2} to c_{iN}.

    12.1.1. Call the current condition c_{ij}. Compare u_{iI}^* to u_{ij}^*.

    12.1.1.1. If they intersect between F_B^i and F_B^{i+1}, identify the smallest such intersection and insert it into the list of left-endpoints as F_B^{i+1}.

    12.1.1.2. Set the incumbent to $\max\{u_{iI}^*, u_{ij}^*\}$, evaluated between F_B^i and the possibly updated value of F_B^{i+1}, noting these two functions do not intersect anywhere in this interval.

    12.2. Record the optimal Red utility and response, given any Blue strategy $F_B^i \leq F_B \leq F_B^{i+1}$, as $u_{R,i}^* = u_{iI}^*$ and $f_{R,i}^* = f_{iI}^*$. Also record the pair of conditions it's known to lie between, denoted $c_{R,i}^{LB}$ and $c_{R,i}^{UB}$.

13. Loop through the list of left-endpoints identified in steps 10 through 12.

13.1. Call the current F_B value F_B^i. If the optimal response (and utility) is the same in $F_B^i \leq F_B \leq F_B^{i+1}$ and $F_B^{i+1} \leq F_B \leq F_B^{i+2}$, and is constrained by the same lower- and upper-bounds, then F_B^{i+1} can be removed from the list. That is, if $f_{R,i}^* = f_{R,i+1}^*$, $c_{R,i}^{LB} = c_{R,i+1}^{LB}$, and $c_{R,i}^{UB} = c_{R,i+1}^{UB}$, then remove F_B^{i+1}.

14. Steps 10 through 13 have determined Red's optimal response, given F_B is within a specified interval. Repeat these steps with F_R and F_B inverted to identify Blue's optimal response function, given F_R.

6 CONCLUSION

Game theory has clearly captured the attention of policy makers, but in its truest form as the mathematical study of strategic interaction, has made limited headway in influencing their decisions. It's clearly made a worthy contribution in influencing the lexicon of qualitative analysis and, on occasion, identified a directional relationship governing human affairs not otherwise obvious. The actual outputs of quantitative models, however, have generally not been used in policy development. This is likely due to: (i) the use of incredibly simplistic and small models; and (ii) a general resistance to mathematical modeling by decision makers who, correctly, recognize the world as more complex than any model. This dissertation sought to show games can have a more direct influence on policy by integrating sophisticated models and techniques into analysis. Point (i) was responded to by dropping traditional assumptions requiring low-level polynomials and minimal constraints in favor of more complex problems, and using efficient algorithms to analyze the corresponding games. Point (ii) was responded to by emphasizing stochasticity and robustness, in the hopes these approaches would pacify decision makers' worries of using the "wrong" model.

Even with stochastic and/or robust approaches, qualitative decision-making will still make the ultimate verdict on a policy decision. It may be that a particular decision maker only wants a model to inform him or her if there's a directional relationship at play in the environment. The analysis conducted here is still valuable in this case: using a more advanced model than a conventional toy model can refute (or confirm) a directional relationship proposed by the latter. In other words, it seems archaic to continue relying on models with simplistic assumptions, even when one's aims are limited. Operations research has produced advanced techniques outside of the game theoretic context that can and ought to be applied to games.

This dissertation provided three examples using advanced operations research techniques in complex games, and the analysis was thus limited to a small set of techniques available. Promising avenues of follow-on research were noted in the conclusions of Chapters 3, 4, and 5, and all essentially centered around expanding the methods used here to bigger and more complex problems. Additional comments on lines of future research would be extraneous, but it's worth emphasizing that the problems analyzed in this dissertation don't come near the level of size and complexity of the most challenging problems policy makers face. For instance, the U.S. Department of Defense routinely uses large-scale simulation models, where single runs to produce one realization of utility may take hours. Contrast this to the problem of Chapter 4, where evaluating the proximity to equilibrium took roughly 3 hours for a model with 60 fisheries, but evaluating a player's utility for a given pair of strategies was virtually instantaneous. Problems as large as these will likely always exceed game theory's ability to analyze them in isolation. In these cases, modelers must consider whether the problem can be decomposed into constituent parts amenable to game-theoretic analysis, and if so, what interactions are at play affecting the conclusions of those individual analyses.

REFERENCES

Acemoglu, Daron, Michael Golosov, Aleh Tsyvinski, and Pierre Yared. 2012. "A Dynamic Theory of Resource Wars." *The Quarterly Journal of Economics* 127 (1): 283–331.

Aghassi, Michele, and Dimitris Bertsimas. 2006. "Robust Game Theory." *Mathematical Programming* 107 (1–2): 231–73. https://doi.org/10.1007/s10107-005-0686-0.

Banks, David L., Jesus Rios, and David Rios Insua. 2016. *Adversarial Risk Analysis*. CRC Press.

Bellhouse, David R., and Nicolas Fillion. 2015. "Le Her and Other Problems in Probability Discussed by Bernoulli, Montmort and Waldegrave." *Statistical Science* 30 (1). https://doi.org/10.1214/14-STS469.

Berman, Eli, Jacob N Shapiro, and Joseph H Felter. 2011. "Can Hearts and Minds Be Bought? The Economics of Counterinsurgency in Iraq." *Journal of Political Economy*, 119 (4): 766-819.

Bielza, Concha, Peter Müller, and David Ríos Insua. 1999. "Decision Analysis by Augmented Probability Simulation." *Management Science* 45 (7): 995–1007. https://doi.org/10.1287/mnsc.45.7.995.

Bier, Vicki, Santiago Oliveros, and Larry Samuelson. 2007. "Choosing What to Protect: Strategic Defensive Allocation against an Unknown Attacker." *Journal of Public*

Economic Theory 9 (4): 563–87. https://doi.org/10.1111/j.1467-9779.2007.00320.x.

Bonferroni, Carlo E. 1936. “Teoria Statistica Delle Classi e Calcolo Delle Probabilita.” *Pubbl. d. R. Ist. Super. Di Sci. Econom. d Commerciali Di Firenze* 8: 1–62.

Borsky, Stefan, and Paul A Raschky. 2011. “A Spatial Econometric Analysis of Compliance with an International Environmental Agreement on Open Access Resources.” *Working Papers in Economics and Statistics*.

Brown, Gerald G., and Louis Anthony Tony Cox Jr. 2011. “How Probabilistic Risk Assessment Can Mislead Terrorism Risk Analysts.” *Risk Analysis* 31 (2): 196–204. https://doi.org/10.1111/j.1539-6924.2010.01492.x.

Brown, Matthew, William B. Haskell, and Milind Tambe. 2014. “Addressing Scalability and Robustness in Security Games with Multiple Boundedly Rational Adversaries.” In *Decision and Game Theory for Security*, edited by Radha Poovendran and Walid Saad, 8840:23–42. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-12601-2_2.

Buer, Tobias, Jörg Homberger, and Hermann Gehring. 2013. “A Collaborative Ant Colony Metaheuristic for Distributed Multi-Level Uncapacitated Lot-Sizing.” *International Journal of Production Research* 51 (17): 5253–70. https://doi.org/10.1080/00207543.2013.802822.

Buszynski, Leszek. 2012. “The South China Sea: Oil, Maritime Claims, and U.S.–China Strategic Rivalry.” *The Washington Quarterly* 35 (2): 139–56. https://doi.org/10.1080/0163660X.2012.666495.

Camerer, Colin. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. The Roundtable Series in Behavioral Economics. New York, N.Y. : Princeton, N.J: Russell Sage Foundation ; Princeton University Press.

Chen, Chun-Hung, and Loo Hay Lee. 2011. *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*. System Engineering and Operations Research, v. 1. Singapore ; Hackensack, NJ: World Scientific.

Chen, Chun-Hung, Jianwu Lin, Enver Yucesan, and Stephen E. Chick. 2000. “Simulation Budget Allocation for Further Enhancing the Efficiency of Ordinal Optimization.” *Discrete Event Dynamic Systems: Theory and Applications* 10: 251–70.

Chen, Weiwei, Siyang Gao, Chun-Hung Chen, and Leyuan Shi. 2014. “An Optimal Sample Allocation Strategy for Partition-Based Random Search.” *IEEE Transactions on Automation Science and Engineering* 11 (1): 177–86. https://doi.org/10.1109/TASE.2013.2251881.

Chick, Stephen E., Jürgen Branke, and Christian Schmidt. 2010. “Sequential Sampling to Myopically Maximize the Expected Value of Information.” *INFORMS Journal on Computing* 22 (1): 71–80. https://doi.org/10.1287/ijoc.1090.0327.

China Power Team. 2020. “Are Maritime Law Enforcement Forces Destabilizing Asia?” *China Power*, April 3, 2020. https://chinapower.csis.org/maritime-forces-destabilizing-asia/.

Chwe, Michael Suk-Young. 2000. “Communication and Coordination in Social Networks.” *Review of Economic Studies* 67 (1): 1–16. https://doi.org/10.1111/1467-937X.00118.

Clark, Colin Whitcomb. 2006. *The Worldwide Crisis in Fisheries: Economic Models and Human Behavior*. Cambridge; New York: Cambridge University Press.

Crespi, Giovanni Paolo, Davide Radi, and Matteo Rocca. 2020. “Insights on the Theory of Robust Games.” *ArXiv Preprint ArXiv:2002.00225*.

De Mesquita, Bruce Bueno. 1997. “A Decision Making Model: Its Structure and Form.” *International Interactions* 23 (3–4): 235–66. https://doi.org/10.1080/03050629708434909.

De Mesquita, Ethan Bueno. 2005. “Conciliation, Counterterrorism, and Patterns of Terrorist Violence.” *International Organization* 59 (1): 145–76.

———. 2007. “Politics and the Suboptimal Provision of Counterterror.” *International Organization* 61 (01). https://doi.org/10.1017/S0020818307070087.

———. 2010. “Regime Change and Revolutionary Entrepreneurs.” *American Political Science Review* 104 (03): 446–66. https://doi.org/10.1017/S0003055410000274.

Edmond, Chris. 2013. “Information Manipulation, Coordination, and Regime Change.” *The Review of Economic Studies* 80 (4): 1422–58. https://doi.org/10.1093/restud/rdt020.

Erickson, Andrew S. 2018. “Numbers Matter: China’s Three ‘Navies’ Each Have the World’s Most Ships.” *National Interest* 26.

Erickson, Andrew S., and Conor M. Kennedy. 2016. “China’s Maritime Militia.” *CNA Corporation* 7: 1–28.

Fan, Weiwei, L. Jeff Hong, and Barry L. Nelson. 2016. “Indifference-Zone-Free Selection of the Best.” *Operations Research* 64 (6): 1499–1514. https://doi.org/10.1287/opre.2016.1530.

Fang, Fei, Peter Stone, and Milind Tambe. 2015. “When Security Games Go Green: Designing Defender Strategies to Prevent Poaching and Illegal Fishing.” In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2589–95.

Fischer, Ronald D., and Leonard J. Mirman. 1996. “The Compleat Fish Wars: Biological and Dynamic Interactions.” *Journal of Environmental Economics and Management* 30 (1): 34–42. https://doi.org/10.1006/jeem.1996.0003.

Glynn, P., and S. Juneja. 2004. “A Large Deviations Perspective on Ordinal Optimization.” In *Proceedings of the 2004 Winter Simulation Conference, 2004.*, 1:565–73. Washington, D.C.: IEEE. https://doi.org/10.1109/WSC.2004.1371364.

Golman, Russell, and Scott E. Page. 2009. “General Blotto: Games of Allocative Strategic Mismatch.” *Public Choice* 138 (3–4): 279–99. https://doi.org/10.1007/s11127-008-9359-x.

Grønbæk, Lone, Marko Lindroos, Gordon R Munro, and Pedro Pintassilgo. 2020. *Game Theory and Fisheries Management: Theory and Applications*. Cham: Springer. https://public.ebookcentral.proquest.com/choice/publicfullrecord.aspx?p=6111258.

Gross, Oliver, and Robert Wagner. 1950. “A Continuous Colonel Blotto Game.” Santa Monica, CA: RAND Project Air Force.

Hanley, Jr., John T. 2018. “Using Game Theory to Identify Opportunities for Deception During Course of Action Analysis and Comparison.” Presented at the 86th Military Operations Research Society Symposium, Monterey, CA, June 21.

Harrison, Glenn W., and E. Elisabet Rutström. 1991. “Trade Wars, Trade Negotiations and Applied Game Theory.” *The Economic Journal* 101 (406): 420–35.

Harsanyi, John C. 1967. “Games with Incomplete Information Played by Bayesian Players, Part I - III.” *Management Science* 14.

Harsanyi, John C., Joseph B. Kadane, and Patrick D. Larkey. 1982. “Subjective Probability and the Theory of Games: Comments, Reply, and Rejoinder between Kadane, Larkey, and Harsanyi.” *Management Science* 28 (2): 120–24.

Hastie, Trevor, Robert Tibshirani, and J. H. Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd ed. Springer Series in Statistics. New York, NY: Springer.

Hsiao, Amanda. 2020. “Opportunities for Fisheries Enforcement Cooperation in the South China Sea.” *Marine Policy* 121: 103569. https://doi.org/10.1016/j.marpol.2019.103569.

Isaacs, Rufus. 1999. *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. Mineola, N.Y: Dover Publications.

Jones, Donald R. 2001. “A Taxonomy of Global Optimization Methods Based on Response Surfaces.” *Ournal of Global Optimization* 21 (4): 345–83.

Jovanovic, Boyan, and Robert W. Rosenthal. 1988. “Anonymous Sequential Games.” *Journal of Mathematical Economics* 17 (1): 77–87.

Kadane, Joseph B., and Patrick D. Larkey. 1982. “Subjective Probability and the Theory of Games.” *Management Science* 28 (2): 113–20.

Kadane, Joseph B., Nicholas Polson, David L. Banks, Francesca Petralia, and Shouqiang Wang. 2011. “Adversarial Risk Analysis: Borel Games. Discussions and Rejoinder between Kadane, Polson, and Banks, et Al.” *Applied Stochastic Models in Business and Industry* 27 (2): 87–94.

Kardes, Erim. 2005. “Robust Stochastic Games and Applications to Counter-Terrorism Strategies.” *CREATE Report*.

King, Dennis M., and Jon G. Sutinen. 2010. “Rational Noncompliance and the Liquidation of Northeast Groundfish Resources.” *Marine Policy* 34 (1): 7–21. https://doi.org/10.1016/j.marpol.2009.04.023.

Kiss, Hubert Janos, Ismael Rodríguez-Lara, and Alfonso Rosa-García. 2017. “Overthrowing the Dictator: A Game-Theoretic Approach to Revolutions and Media.” *Social Choice and Welfare* 49 (2): 329–55. https://doi.org/10.1007/s00355-017-1067-3.

Koller, Daphne, and Brian Milch. 2003. “Multi-Agent Influence Diagrams for Representing and Solving Games.” *Games and Economic Behavior* 45 (1): 181–221. https://doi.org/10.1016/S0899-8256(02)00544-4.

Kovenock, Dan, and Brian Roberson. 2012a. *Conflicts with Multiple Battlefields*. Oxford University Press. https://doi.org/10.1093/oxfordhb/9780195392777.013.0021.

———. 2012b. "Coalitional Colonel Blotto Games with Application to the Economics of Alliances." *Journal of Public Economic Theory* 14 (4): 653–76. https://doi.org/10.1111/j.1467-9779.2012.01556.x.

Kuperan, K., and Jon G. Sutinen. 1998. "Blue Water Crime: Deterrence, Legitimacy, and Compliance in Fisheries." *Law & Society Review* 32 (2): 309–38. https://doi.org/10.2307/827765.

Lapan, Harvey E., and Todd Sandler. 1988. "To Bargain or Not To Bargain: That Is The Question." *The American Economic Review* 78 (2,): 16–21.

Liang, Xiannuan, and Yang Xiao. 2013. "Game Theory for Network Security." *IEEE Communications Surveys & Tutorials* 15 (1): 472–86. https://doi.org/10.1109/SURV.2012.062612.00056.

Lin, Kyle, and Jeffrey Dayton. 2011. "Game-Theoretic Models for Jamming Radio-Controlled Improvised Explosive Devices." *Military Operations Research* 16 (3): 5–13. https://doi.org/10.5711/1082598316305.

Long, Le Kim, and Ola Flaaten. 2011. "A Stackelberg Analysis of the Potential for Cooperation in Straddling Stock Fisheries." *Marine Resource Economics* 26 (2): 119–39. https://doi.org/10.5950/0738-1360-26.2.119.

Mangin, Tracey, Christopher Costello, James Anderson, Ragnar Arnason, Matthew Elliott, Steve D. Gaines, Ray Hilborn, Emily Peterson, and Rashid Sumaila. 2018. "Are Fishery Management Upgrades Worth the Cost?" Edited by Athanassios C.

Tsikliras. *PLOS ONE* 13 (9): e0204258. https://doi.org/10.1371/journal.pone.0204258.

McKelvey, Richard D., and Thomas R. Palfrey. 1995. “Quantal Response Equilibria for Normal Form Games.” *Games and Economic Behavior* 10 (1): 6–38.

McLay, Laura, Casey Rothschild, and Seth Guikema. 2012. “Robust Adversarial Risk Analysis: A Level-k Approach.” *Decision Analysis* 9 (1): 41–54. https://doi.org/10.1287/deca.1110.0221.

Miller, Steve, and Bruno Nkuiya. 2016. “Coalition Formation in Fisheries with Potential Regime Shift.” *Journal of Environmental Economics and Management* 79: 189–207.

Nash, John. 1951. “Non-Cooperative Games,” September, 11.

———. 1953. “Two-Person Cooperative Games.” *Econometrica* 21: 128–40.

Neumann, John von. 1928. “Zur Theorie Der Gesellschaftsspiele.” *Mathematical Annals* 100.

Neumann, John von, and Oskar Morgenstern. 1944. *Theory of Games and Economic Behaviour*. Princeton University Press.

Nguyen, Thanh H, Rong Yang, Amos Azaria, Sarit Kraus, and Milind Tambe. 2013. “Analyzing the Effectiveness of Adversary Modeling in Security Games.” In , 7.

Nielsen, Jesper Raakjær, and Christoph Mathiesen. 2003. “Important Factors Influencing Rule Compliance in Fisheries Lessons from Denmark.” *Marine Policy* 27 (5): 409–16.

Nikoofal, Mohammad E., and Jun Zhuang. 2012. “Robust Allocation of a Defensive Budget Considering an Attacker’s Private Information: Robust Allocation of a Defensive Budget.” *Risk Analysis* 32 (5): 930–43. https://doi.org/10.1111/j.1539-6924.2011.01702.x.

Nisan, Noam, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, eds. 2007. *Algorithmic Game Theory*. Cambridge ; New York: Cambridge University Press.

Park, Young-kil. 2020. “The Role of Fishing Disputes in China–South Korea Relations.” The National Bureau of Asian Research (NBR). https://map.nbr.org/2020/04/the-role-of-fishing-disputes-in-china-south-korea-relations/.

Peng, Yijie, Chun-Hung Chen, Michael C. Fu, and Jian-Qiang Hu. 2016. “Dynamic Sampling Allocation and Design Selection.” *INFORMS Journal on Computing* 28 (2): 195–208. https://doi.org/10.1287/ijoc.2015.0673.

Perry, Michael M. 2020. “Cooperative Maritime Law Enforcement and Overfishing in the South China Sea.” *Center for International Maritime Security*, April 6, 2020. http://cimsec.org/cooperative-maritime-law-enforcement-and-overfishing-in-the-south-china-sea/43227.

Petrossian, Gohar A. 2015. “Preventing Illegal, Unreported and Unregulated (IUU) Fishing: A Situational Approach.” *Biological Conservation* 189: 39–48. https://doi.org/10.1016/j.biocon.2014.09.005.

———. 2019. *The Last Fish Swimming: The Global Crime of Illegal Fishing*. Global Crime and Justice. Santa Barbara, California: Praeger, an imprint of ABC-CLIO, LLC.

Pitcher, Tony, Daniela Kalikoski, and Ganapathiraju Pramod. 2006. “Evaluations of Compliance with the FAO (UN) Code of Conduct for Responsible Fisheries.” *Fisheries Centre Research Reports, University of British Columbia* 14 (2).

Powell, Michael J.D. 1998. “Direct Search Algorithms for Optimization Calculations.” *Acta Numerica*, 287–336.

Rios Insua, David, Jesus Rios, and David L. Banks. 2009. “Adversarial Risk Analysis.” *Journal of the American Statistical Association* 104 (486): 841–54. https://doi.org/10.1198/jasa.2009.0155.

Rios, Jesus, and David Rios Insua. 2012. “Adversarial Risk Analysis for Counterterrorism Modeling.” *Risk Analysis* 32 (5): 894–915. https://doi.org/10.1111/j.1539-6924.2011.01713.x.

Rosenberg, David. 2005. “Managing the Resources of the China Seas: China’s Bilateral Fisheries Agreements with Japan, South Korea, and Vietnam.” *The Asia-Pacific Journal* 3 (6).

Rubinstein, Ariel. 2012. *Economic Fables*. Cambridge: Open Book Publ. https://books.openbookpublishers.com/10.11647/obp.0020/toc.html.

———. n.d. The best books on Game Theory: recommended by Ariel Rubinstein. https://fivebooks.com/best-books/ariel-rubinstein-on-game-theory/.

Schelling, Thomas C. 1980. *The Strategy of Conflict: With a New Preface*. Nachdr. d. Ausg. 1980. Cambridge, Mass.: Harvard Univ. Press.

Sevillano, Juan Carlos, David Rios Insua, and Jesus Rios. 2012. "Adversarial Risk Analysis: The Somali Pirates Case." *Decision Analysis* 9 (2): 86–95. https://doi.org/10.1287/deca.1110.0225.

Shi, Leyuan, and Sigurdur Ólafsson. 2000. "Nested Partitions Method for Global Optimization." *Operations Research* 48 (3): 390–407. https://doi.org/10.1287/opre.48.3.390.12436.

Smith, J. Maynard, and George R. Price. 1973. "The Logic of Animal Conflict." *Nature* 246 (5427): 15–18.

Sorokin, Gerald L. 1994. "Alliance Formation and General Deterrence: A Game-Theoretic Model and the Case of Israel." *Journal of Conflict Resolution* 38 (2): 298–325.

"South China Sea Threatened by 'a Series of Catastrophes.'" 2019. *PBS NewsHour*. https://www.pbs.org/newshour/show/south-china-sea-threatened-by-a-series-of-catastrophes.

Stearns, Scott. 2012. "Challenging Beijing in the South China Sea." Voice of America. *State of Affairs* (blog). July 31, 2012. https://blogs.voanews.com/state-department-news/2012/07/31/challenging-beijing-in-the-south-china-sea/.

Sureka, Ashish, and Peter R. Wurman. 2005. "Applying Metaheuristic Techniques to Search the Space of Bidding Strategies in Combinatorial Auctions." In *Proceedings of the 2005 Conference on Genetic and Evolutionary Computation - GECCO '05*, 2097. Washington DC, USA: ACM Press. https://doi.org/10.1145/1068009.1068356.

Wang, Chen, and Vicki M. Bier. 2011. “Target-Hardening Decisions Based on Uncertain Multiattribute Terrorist Utility.” *Decision Analysis* 8 (4): 286–302. https://doi.org/10.1287/deca.1110.0218.

Wang, Shouqiang, and David L. Banks. 2011. “Network Routing for Insurgency: An Adversarial Risk Analysis Framework.” *Naval Research Logistics (NRL)* 58 (6): 595–607. https://doi.org/10.1002/nav.20469.

Wilson, Edwin B. 1927. “Probable Inference, the Law of Succession, and Statistical Inference.” *Journal of the American Statistical Association* 22 (158): 209–12.

Xu, Jie, Edward Huang, Chun-Hung Chen, and Loo Hay Lee. 2015. “Simulation Optimization: A Review and Exploration in the New Era of Cloud Computing and Big Data.” *Asia-Pacific Journal of Operational Research* 32 (03): 1550019. https://doi.org/10.1142/S0217595915500190.

Xu, Jie, Si Zhang, Edward Huang, Chun-Hung Chen, Loo Hay Lee, and Nurcin Celik. 2016. “Multi-FIdelity Optimization with Ordinal Transformation and Optimal Sampling.” *Asia-Pacific Journal of Operational Research* 33 (3): 26.

Yonhap News. 2020. “S. Korea, China Agree to Cut Fishing in EEZs,” November 6, 2020. https://en.yna.co.kr/view/AEN20201106010700320.

Yoshihara, Toshi. 2012. “China’s Vision of Its Seascape: The First Island Chain and Chinese Seapower: China’s Vision of Its Seascape.” *Asian Politics & Policy* 4 (3): 293–314. https://doi.org/10.1111/j.1943-0787.2012.01349.x.

Zhang, Hongzhou. 2016. “Chinese Fishermen in Disputed Waters: Not Quite a ‘People’s War.’” *Marine Policy* 68: 65–73. https://doi.org/10.1016/j.marpol.2016.02.018.

Zhang, Hongzhou, and Sam Bateman. 2017. “Fishing Militia, the Securitization of Fishery and the South China Sea Dispute.” *Contemporary Southeast Asia* 39 (2): 288–314. https://doi.org/10.1355/cs39-2b.

Zhuang, Jun, and Vicki M. Bier. 2007. “Balancing Terrorism and Natural Disasters—Defensive Strategy with Endogenous Attacker Effort.” *Operations Research* 55 (5): 976–91. https://doi.org/10.1287/opre.1070.0434.

BIOGRAPHY

Michael Macgregor Perry was born in 1985 and joined the U.S. Navy in 2009.