

Air Force Institute of Technology

AFIT Scholar

Theses and Dissertations

Student Graduate Works

6-2019

Persuasion, Political Warfare, and Deterrence: Behavioral and Behaviorally Robust Models

William N. Caballero

Follow this and additional works at: <https://scholar.afit.edu/etd>



Part of the [Behavioral Economics Commons](#), and the [Operational Research Commons](#)

Recommended Citation

Caballero, William N., "Persuasion, Political Warfare, and Deterrence: Behavioral and Behaviorally Robust Models" (2019). *Theses and Dissertations*. 4467.

<https://scholar.afit.edu/etd/4467>

This Dissertation is brought to you for free and open access by the Student Graduate Works at AFIT Scholar. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of AFIT Scholar. For more information, please contact richard.mansfield@afit.edu.



**PERSUASION, POLITICAL WARFARE, AND
DETERRENCE: BEHAVIORAL AND
BEHAVIORALLY ROBUST MODELS**

DISSERTATION

William N. Caballero, Capt, USAF

AFIT-ENS-DS-19-J-022

**DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY**

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

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

AFIT-ENS-DS-19-J-022

PERSUASION, POLITICAL WARFARE, AND DETERRENCE: BEHAVIORAL
AND BEHAVIORALLY ROBUST MODELS

DISSERTATION

Presented to the Faculty
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
in Partial Fulfillment of the Requirements for the
Degree of Doctor of Philosophy

William N. Caballero, B.S., M.S.
Capt, USAF

13 June 2019

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

AFIT-ENS-DS-19-J-022

PERSUASION, POLITICAL WARFARE, AND DETERRENCE: BEHAVIORAL
AND BEHAVIORALLY ROBUST MODELS

DISSERTATION

William N. Caballero, B.S., M.S.
Capt, USAF

Committee Membership:

Dr. Brian J. Lunday, PhD
Chair

Dr. Richard F. Deckro, DBA
Member

Dr. Meir N. Pachter, PhD
Member

Abstract

This dissertation examines game theory models in the context of persuasion and competition wherein decisionmakers are boundedly rational by considering two complementary threads of research. The first thread of research pertains to offensive and preemptively defensive behavioral models of influence, and the second thread of research discussed herein pertains to behavioral and behaviorally robust approaches to deterrence and other military operations via the utilization of behavioral game theory and mathematical programming under uncertainty.

Persuasion is a fundamental element of human interaction applied to both individuals and populations. Although persuasion is a well-studied, interdisciplinary field of research, this work advances its prescriptive, quantitative characterization and future use. That is, this research complements the qualitative psychological literature with respect to the processing of persuasive messages by developing an offensive influence campaign design framework. We adapt the classic Decision Analysis problem to a bilevel mathematical program, wherein a persuader has the opportunity to affect the environment prior to the decisionmaker's choice. Thereby, we define a new class of problems for modeling persuasion. Utilizing Cumulative Prospect Theory as a descriptive framework of choice, we transform the persuasion problem to a single level mathematical programming formulation, adaptable to conditions of either risk or uncertainty. These generalized models allow for the malleability of prospects as well as Cumulative Prospect Theory parameters through persuasion update functions. We detail the literature that supports the quantification of such effects which, in turn, establishes that such update functions can be realized. Finally, the efficacy of the model is illustrated through three use cases under varying conditions of risk or un-

certainty: the establishment of insurance policies, the construction of a legal defense, and the development of a public pension program.

However, in an influence setting, it may be the case that multiple actors compete over a population's decisions. This is especially true in modern international relations. As such, this work presents two new game theoretic frameworks, denoted as *prospect games* and *regulated prospect games*, to inform defensive policy against these threats. These frameworks respectively model (a) the interactions of competing entities influencing a populace and (b) the preemptive actions of a regulating agent to alter such a framework. Prospect games and regulated prospect games are designed to be adaptable, depending on the assumed nature of persuaders' interactions and their rationality. The contributions herein are a modeling framework for competitive influence operations under a common set of assumptions, model variants that respectively correspond to scenario-specific modifications of selected assumptions, the illustration of practical solution methods for the suite of models, and a demonstration on a representative scenario with the ultimate goal of providing a quantifiable, tractable, and rigorous framework upon which national policies defending against competitive influence can be identified.

Moreover, even when acting in isolation, the central task of an influencing entity is confounded by uncertainty of either a structural or parametric form. The research herein also sets forth a modeling framework to identify robust influence strategies under such uncertain conditions. Furthermore, the utility of this framework and its proper parameterization are illustrated via an application to the contemporary, global problem of the Afghan opium trade. Utilizing a variety of open source data, we demonstrate how counternarcotic policy can be informed using a quantitative method that embraces the bounded rationality of the economy's decisionmakers and the government's uncertainty regarding the degree of this deviation from rationality.

In this manner, we provide a new framework from which robust influence decisions can be made under realistic information conditions, and elucidate how it can be used to inform real-world policy.

Related to the second thread of research, since Thomas Schelling published *The Strategy of Conflict* (1960), the study of game theory and international relations have been closely linked. Developments in the former often trigger analytical changes in the latter, as evidenced by the recent behavioral and psychological focus among some international relations and defense economics scholars. Despite this connection, decisions regarding military operations have rarely been influenced by game theoretic analysis, a fact often attributed to standard game theory's normative nature. Therefore, this research applies selected behavioral game theoretic solution techniques to classical interstate conflict games, demonstrating their utility to inform the planning of military operations. By reexamining classic Cold War deterrence models and other interstate conflict games, we demonstrate how modern game theoretic techniques based upon agent psychology, as well as the ability of agents to think strategically or learn from past experience, can provide additional insights beyond what can be derived via perfect rationality analysis. These demonstrations illustrate how behaviorally focused methods can incorporate the uncertainty related to human decisionmakers into analysis and highlight the alternative insights a bounded rationality approach can generate for military operations planning.

Under the assumption that a nation's adversaries are boundedly rational, the final contribution defines a methodology for strategy identification. That is, recent advances in behavioral game theory address a persistent criticism of traditional solution concepts that rely upon perfect rationality: equilibrium results are often inconsistent with empirical evidence. For normal-form games, the Cognitive Hierarchy model is a solution concept based upon a sequential reasoning process, yielding accurate char-

acterizations of experimental human game play. However, to derive predictions using Cognitive Hierarchy, an accurate estimate of the average number of reasoning steps players utilize is required. Such information is generally unknown *ex ante* and is a source of uncertainty in practice. Moreover, the Cognitive Hierarchy model, like much research in game theory, analyzes a game holistically. Herein, we adopt a different approach, considering the normal-form game as a decision problem from the perspective of an arbitrary player. Assuming such a player is confronting a set of boundedly rational opponents whose play is characterized by the Cognitive Hierarchy model, we develop a suite of six mathematical programming formulations to maximize the player's minimum payoff, with the appropriate model identifiable under varying levels of information with respect to their opponents' reasoning abilities. By leveraging robust optimization, stochastic programming, and distributionally robust optimization techniques, our set of models yields prescriptive strategies of how a normal-form game should be played. A software package implementing these constructs is developed and applied to illustrative instances, demonstrating how these behaviorally robust strategies vary in accordance with the underlying uncertainty.

To my wife and sons, you are my inspiration and motivation.

Acknowledgements

Without the mentorship of Dr. Brian Lunday, this dissertation would not have been possible. Thank you for allowing me to follow my interests and for guiding this research. I am a better analyst and officer because of your tutelage.

Thank you also to the members of my committee, Dr. Richard Deckro and Dr. Meir Pachter, for your continued support in the creation and improvement of this manuscript.

William N. Caballero

Table of Contents

	Page
Abstract	iv
Acknowledgements	ix
List of Figures	xiii
List of Tables	xv
I. Introduction	1
1.1 Motivation	1
1.2 Research Objectives and Scope	3
1.3 Organization of the Dissertation	5
II. Influence Modeling: Mathematical Programming Representations of Persuasion under Either Risk or Uncertainty	7
2.1 Introduction	8
2.2 Relevant Literature	14
Persuasion	14
Other Models for Decision Making Under Uncertainty	18
2.3 Influence Modeling	19
The Parameter Update Functions	20
Proposed Model under Risk	23
Adaptations from Risk to Uncertainty	33
2.4 Example Applications	36
Influence under Risk: Insurance Policies	38
Influence under Uncertainty with Evidence Strength Metric: Defending a Client	43
Influence under Uncertainty without Evidence Strength Metric: Pension Enrollment	49
2.5 Discussion and Conclusions	53
III. Informing National Security Policy by Modeling Adversarial Inducement and its Governance	56
3.1 Introduction	56
3.2 Prospect Games	60
Defining Prospect Games	63
On the Malleability of Prospects and CPT Parameters	68
Solving Prospect Games	69
3.3 Regulated Prospect Games	71

	Page
Defining Regulated Prospect Games	72
Solving Regulated Prospect Games	80
3.4 Example Application: Deterring Electoral Interference.....	83
Perfectly Rational Persuaders: Nash Equilibrium	88
Perfectly Rational Persuaders: Correlated Equilibrium	92
Boundedly Rational Persuaders: Cognitive Hierarchy.....	94
Uncertainty in the Prospect Game Payoff Structure	97
Discussion	100
3.5 Conclusions.....	101
IV. Robust Influence Modeling under Structural and Parametric Uncertainty: An Afghan Counternarcotics Use Case.....	103
4.1 Introduction	103
4.2 Robust Decision Making and Influence Modeling.....	105
Robust Decision Making.....	108
An RDM Approach to Influence	110
4.3 Case Study: Influencing Landowning Household Crop Choice in Badakhshan Province	115
Seasonal Crop Decision by Badakhshan Landowners.....	117
Influence Actions and their Effects	126
RDM Tailorable Components	129
Analysis and Results	131
Discussion and Limitations	135
4.4 Conclusions.....	137
V. Leveraging Behavioral Game Theory to Inform Military Operations Planning	139
5.1 Introduction	139
5.2 Behavioral Game Theory	143
Cognitive Hierarchy.....	144
Experience Weighted Attraction	147
5.3 Applications to Military Operations Planning	149
Brinkmanship: Nuclear Crisis Game	151
Target Selection: D-Day Game	155
First Strike Decision: Preemptive War Game.....	159
Incorporating Behavioral Uncertainty	163
5.4 Alternative Game Models and BGT Modeling Approaches	166
5.5 Conclusions and Recommendations	169

	Page
VI. Identifying Behaviorally Robust Maximin Strategies for Normal-form Games under Varying Forms of Uncertainty	171
6.1 Introduction	172
6.2 Review of the Cognitive Hierarchy Model	176
6.3 Behaviorally Robust Strategies in Normal Form Games	178
Finite Uncertainty Set over τ	182
Interval Uncertainty Set over τ	186
6.4 Behaviorally Robust Strategies for the Stahl & Wilson Games	197
The Twelve Stahl & Wilson Games	197
6.5 Conclusion	203
VII. Conclusions	205
Appendix A. Expected Utility Influence Model	206
Appendix B. IISE Proceedings: <i>Challenges and Solutions with Exponentiation Constraints using Decision Variables via the BARON Commercial Solver</i>	207
Bibliography	214

List of Figures

Figure		Page
1	A Decision Analysis Visualization of International Conflict	3
2	Simple Influence Example under Risk for Risk-Neutral, EUT Decisionmaker	10
3	Customer Decision Tree for Purchasing Insurance Policy	38
4	Juror Decision Tree: Two Available Verdicts	44
5	Pension Enrollment Decision Tree	50
6	Compressed Pension Enrollment Decision Tree	51
7	Example Decision Tree	59
8	PG with 3 Persuaders each having 3 Actions, and n decisionmakers	61
9	Structure of Regulated Prospect Games	71
10	Deterring Improper Electoral Interference as a Regulated Prospect Game	84
11	Induced Prospect Game with No Sanctions Imposed	88
12	Distributions of Finishes for Optimistic-Optimal Nash Equilibrium RPG	91
13	Distributions of Finishes under Optimal Correlated Equilibrium RPG	94
14	Distributions of Finishes under Optimal Cognitive Hierarchy RPG	96
15	Example Decision Tree between Course of Actions (COAs) 1 and 2	106
16	Badakhshan Province, Afghanistan	117
17	Generic Badakhshan Farmer Decision Tree	121
18	Distribution of Poppy Cultivated Jeribs over all Futures	131

Figure	Page
19	Distribution of Absolute Regret over all Futures 132
20	Cluster Tradeoffs Between s_{16} and Other Strategies 133
21	Empirical CDFs for s_{12} with Best- and Worst-Case Opium Yields 135
22	Player strategies as a function of τ 152
23	Probability of war as a function of τ 153
24	Players strategies over time in 100 simulated EWA runs 154
25	Number of time periods before nuclear exchange 155
26	Probability Allies attack a well-fortified location as a function of τ 157
27	Allies' strategies with $\tau = 0, 0.1, 0.2, \dots, 50$ 158
28	Germans' strategies with $\tau = 0, 0.1, 0.2, \dots, 50$ 158
29	Probability No Attack in Prisoner's Dilemma as a function of τ 160
30	Probability No Attack in Assurance game as a function of τ 161
31	Empirical vs. simulated mean for Prisoner Dilemma fitted EWA 162
32	Row Player Probability Mass Function for Likelihood of Deescalation 164
33	M -step Column Player Expected Payoffs in Stahl & Wilson Game 10 181
34	M -step Column Player Expected Payoffs in Stahl & Wilson Game 1 200
35	M -step Column Player Expected Payoffs in Stahl & Wilson Game 4 200

List of Tables

Table	Page
1	Correspondence of Instance-specific and GPP constraints 41
2	Parameter Values Associated with Insurance Persuasion Program 42
3	Parameter Update Coefficients Associated with Defense Themes 47
4	Correspondence of Instance-specific and GPP constraints 48
5	Parameter Update Coefficients Associated with Pension Persuasion 51
6	Correspondence of Instance-specific and GPP constraints 52
7	Uniform Distribution on voter CPT-parameters, and \hat{x}_{ijk} - and \hat{p}_{ijk} -values 85
8	Regulator Objective Function Coefficient Values 87
9	Election Inference RPG - Optimal Optimistic-Nash Equilibriums 90
10	Election Inference RPG - Optimal Correlated Equilibriums 92
11	Election Inference PG - Cognitive Hierarchy Solution 95
12	Election Inference RPG - Cognitive Hierarchy Solutions 95
13	Solutions of PGs for 10 Additional Scenarios 98
14	Solutions of RPGs for 10 Additional Scenarios 99
15	Estimated Per Jerib Net Income Range 123
16	Decisionmaker Uncertainty Sets 126
17	Actions Included (Y) and Excluded (N) in each Influence Strategy 128
18	Influence Effects on Decision Setting 129

Table	Page
19	Two-player Nuclear Crisis Game 151
20	D-Day Game 156
21	Preemptive War: Prisoner's Dilemma Game 160
22	Preemptive War: Assurance Game 160
23	Allies' Payoff and Uncertainty Information 166
24	Game 10 from Stahl and Wilson (1995) 180
25	Seven information conditions for Stahl & Wilson games 198
26	Prescribed mixed strategies for all Stahl & Wilson games 199
27	Descriptive expected payoff statistics versus 48 Stahl & Wilson players 201

PERSUASION, POLITICAL WARFARE, AND DETERRENCE: BEHAVIORAL
AND BEHAVIORALLY ROBUST MODELS

I. Introduction

1.1 Motivation

The evolution of war's manifestation - its *character* - is evident in history. The development of stirrups enabled Genghis Kahn's armies to advance and apply cavalry tactics to overwhelm their enemies and conquer the steppes of Asia (Saunders, 2001). The advent of the longbow and the use of highly trained archers changed the character of war as exhibited at the Battle of Crecy (Burne, 2016). More recently, beginning in the American Civil War and more fully exhibited during World War I, the development of weapons having greater firepower and higher rates of fire, along with the tactics to use them, marked the ascendancy of the defense over the offense in wars of attrition or exhaustion (United States Military Academy, 2014a,b). In World War II, the character of warfare changed with the use of armored forces, close air support, and well-coordinated combined arms (United States Military Academy, 2015a,b).

However, students of Clausewitz would argue that such changes in war's manifestation do not imply a change in its essence, i.e., the *nature* of war. Although technology and tactics may evolve over time, its nature is unchanging.

Clausewitz (1989) described war as an "instrument of policy" that uses "violence intended to compel our opponent to fulfill our will." Therefore, warfare can be considered to be a type of violent persuasion. An enemy is confronted with multiple courses

of actions, and we try to impel them to choose one in accordance with our preferences. In the contemporary operating environment, other methods beyond physical violence are often utilized to achieve the same objective. The character of warfare is currently undergoing a shift from a kinetic to a narrative focus. Fueled by modern technology, the weaponized use of communication is a critical piece of political warfare (Kennan, 1948; Boot and Doran, 2013; Polyakova and Boyer, 2018), hybrid warfare (United States Government Accountability Office, 2010; Kofman and Rojansky, 2015), influence operations (Larson et al., 2009), and information operations (United States Joint Chiefs of Staff, 2012; Air University, 2018).

Therefore, we contend that key elements of war's underlying nature can be represented within the purposeful manipulation of an adversary's decision tree. Such manipulation may affect how the adversary perceives the underlying uncertainty, values the payoff associated with each outcome, and/or evaluates the available set of prospects. In conventional conflict, examples of how these effects may be achieved respectively include military deception (e.g., Operation Fortitude in WWII), combat engineering (e.g., countermobility obstacles and defensive structures), and shows of force (e.g., Freedom of Navigation operations). If such actions are successful, an adversary acts in accordance with our will, and the underlying desired effect has been achieved.

When considering international competition prior to war, a similar dynamic holds. However, we must now consider manipulating a decision tree via alternative means. A simplification of this dynamic can be seen in Figure 1 wherein the adversary wishes to adopt COA 1, but we desire them to select COA 2. In conventional conflict, these underlying choices are replaced with the respective proxies of "Fight" and "Surrender".

Moreover, it is interesting to note that this type of hybrid warfare does not limit




Figure 1. A Decision Analysis Visualization of International Conflict

itself to affecting the decisions of a nation's leadership. The target of an attack is often a nation's citizenry. Russian aggression has illustrated this numerous times on both the United States and its NATO allies (DHS and FBI, 2016; The Economist, 2019). These tactics have been shown to be sufficiently effective that some experts suggest such acts of political warfare may become the "default mode for conflict in the coming decades" (Cohen and Robinson, 2018).

1.2 Research Objectives and Scope

This research aims to address this evolved character of 21st Century warfare. New technologies must be accompanied by changes in tactics to use them effectively. Political warfare is a reality, and its current application is best characterized as an *art*, because we currently lack the *science* of tools and tactics necessary to best utilize it in either an offensive or defensive setting. This work seeks to make a first step towards developing useful models to advance strategy and tactics in conflicts below the threshold of conventional war.

To ensure the effectiveness of these modeling efforts, we advocate for the use of behavioral theories. Various studies (e.g., Kahneman and Tversky, 1979; Selten, 1998; Gigerenzer and Selten, 2002) have shown human decision making to diverge from perfect rationality and, when building models of persuasion, such limitations

must be considered. Behavioral theories are informed by psychological studies, often taking the form of human subject testing, and have proven themselves effective to describe human behavior.

This dissertation examines game theory models in the context of persuasion and competition wherein decisionmakers are not completely rational by considering two complementary threads of research. The first thread of research pertains to offensive and preemptively defensive behavioral models. Research in this thread makes three notable contributions. First, an offensive modeling framework (i.e., a *persuasion program*) is created to identify how an entity optimally influences a populace to take a desired course of action. In doing so, we summarize the foundational literature upon which future research can build to effectively parameterize such models, and illustrate the utility of the framework over a variety of civil applications. Second, a defensive modeling framework is defined wherein a regulating entity (e.g., a government) takes action to bound the behavior of multiple adversaries simultaneously attempting to persuade a group of decisionmakers. Third, an offensive influence modeling framework under conditions of ambiguity is developed in accordance with historical information limitations, and we demonstrate how it can be used to select a robust course of action on a specific, data-driven use case.

The second thread of research pertains to behavioral and behaviorally robust approaches to deterrence and other military operations. Research in this thread makes two notable contributions. First, we demonstrate the alternative insights behavioral game theory generates for the analysis of classic deterrence games, and explicate the rich analysis generated from its combined use with standard equilibrium models. Second, we define behaviorally robust models for an agent to use in a normal form game under varying forms of uncertainty in order to inform deterrence policy decisions.

Both of these research threads pertain to some form of influence; the first considers

influence at a more tactical level (e.g., a targeted information operation) whereas the second is more strategic (e.g., deterring hostile enemy aggression). Both research threads rely on descriptive theories and explore solution methods when the adversary's behavior is uncertain; however, the first thread focuses on behavioral theories of choice and the second on behavioral game theory.

Finally, we note that influence is a central component of a variety of civilian activities as well (e.g., marketing, criminal justice, politics). As such, the models developed herein are described in general terms and their relevance to national security is elucidated via representative use cases.

1.3 Organization of the Dissertation

Chapters II through IV of this dissertation pertain to the development of offensive and defensive behavioral influence models. Chapter II defines the *general persuasion program* enabling an influencing entity to optimally engage a population under conditions of either risk or ambiguity. This generalized model utilizes tools from Cumulative Prospect Theory, Support Theory, and nonlinear optimization to identify optimal offensive behavior in an influence operation. Chapter III builds upon the general persuasion program by modeling additional influencing entities as well as a preemptive, regulating agent. Using non-linear optimization techniques, a suite of models are developed for the regulating agent to bound the behavior of the influencing entities under varying rationality assumptions and game theoretic solution concepts. In Chapter IV, we extend the general persuasion program presented in Chapter II while retaining its original persuader-populace structure. We focus on parametric and structural uncertainty associated with the populace's behavior, and the identification of behaviorally robust strategies given these knowledge limitations. The efficacy of this methodology for real-world policy development is then illustrated on an Afghan

counternarcotic use case informed by open-source UN, American, Afghan and NGO data.

Chapters V and VI concern behavioral and behaviorally robust approaches to deterrence. In Chapter V we illustrate how the Cognitive Hierarchy and Experience Weighted Attraction algorithms can be utilized to inform military operations planning. This topic is extended in Chapter VI, wherein we develop a suite of mathematical programs for the selection of robust strategies given this underlying behavioral uncertainty.

Chapter VII provides concluding remarks and the two appendices provide relevant, supplemental material. The first appendix illustrates how the general persuasion programs of Chapter II can be extended when decisionmakers are assumed to maximize expected utility. The second appendix documents an ancillary research contribution regarding the identification of erroneous reports of instance infeasibility when solving power program problems with a leading commercial solver for global optimization (i.e., BARON), as well as the development and testing of theoretically-grounded modeling techniques to reduce their occurrence, thereby enhancing solver efficacy. Such a computational challenge was identified (and resolved favorably, in most instances tested) when solving a simplified form of the persuasion program presented in Chapter II.

II. Influence Modeling: Mathematical Programming Representations of Persuasion under Either Risk or Uncertainty

Abstract

Persuasion is a fundamental element of human interaction applied to both individuals and populations. Although persuasion is a well-studied, interdisciplinary field of research, this work advances its prescriptive, quantitative characterization and future use. That is, this research complements the qualitative psychological literature with respect to the processing of persuasive messages by developing an influence campaign design framework. We adapt the classic Decision Analysis problem to a bilevel mathematical program, wherein a persuader has the opportunity to affect the environment prior to the decisionmaker's choice. Thereby, we define a new class of problems for modeling persuasion. Utilizing Cumulative Prospect Theory as a descriptive framework of choice, we transform the persuasion problem to a single level mathematical programming formulation, adaptable to conditions of either risk or uncertainty. These generalized models allow for the malleability of prospects as well as Cumulative Prospect Theory parameters through persuasion update functions. We detail the literature that supports the quantification of such effects which, in turn, establishes that such update functions can be realized. Finally, the efficacy of the model is illustrated through three use cases under varying conditions of risk or uncertainty: the establishment of insurance policies, the construction of a legal defense, and the development of a public pension program.

2.1 Introduction

Persuasion is ever-present in human society. Television advertisements persuade consumers to buy a product. Parents convince children to eat their vegetables. Laws outlining punishment persuade citizens to abide by certain behaviors. A suitor persuades their beloved to marry, and a government attempts to influence other countries in the realm of international relations. Such ubiquitous activities are well-studied in a variety of disciplines but are not addressed frequently in a prescriptive, quantitative manner. Therefore, we formalize herein a new class of decision problems which can mathematically model persuasion and its influence on others, in either situations of risk when outcome probabilities are known, or uncertainty when they are unknown.

Many persuasion scenarios can be modeled with a framework similar to that used in the discipline of Decision Analysis. Using Prescriptive Decision Analysis when outcomes of a given choice are not deterministic, an advisor often represents problems as decision trees to help a client take the best action in accordance with the axioms of Expected Utility Theory (EUT). Such a methodology is applicable in the context of risk or uncertainty (e.g., casino games or stock prices, respectively). The approach is identical in either setting with the exception that uncertainty requires a subjective assessment of outcome probabilities. The generic Decision Analysis problem of an individual choosing between $|J|$ options (i.e., choosing among a set of prospects, J) can be represented via a mathematical programming formulation as follows.

$$\begin{aligned} & \max_{\zeta} \sum_{j \in J} \zeta_j EU(j) \\ & \text{subject to } \sum_{j \in J} \zeta_j = 1, \end{aligned} \tag{1a}$$

$$\zeta_j \in \{0, 1\}, \forall j \in J. \tag{1b}$$

The value $EU(j) = \sum_{k \in K_j} p_{jk} u_{jk}$, where K_j are all possible outcomes that can occur from a selection of prospect j , each occurring with non-negative probability p_{jk} and having utility u_{jk} units to the decisionmaker. At optimality, the prospect corresponding to $\zeta_j^* = 1$ is the option that will maximize the decisionmaker's expected utility.

We now reexamine the Prescriptive Decision Analysis problem by assuming the advisor is no longer altruistic but self-interested with a distinct preference on what alternative the decisionmaker selects. Consider a perfectly rational, risk-neutral decisionmaker selecting between two risky prospects as illustrated by the decision tree shown in Figure 2, and a persuader with three available actions (i.e., a_1, a_2 , and a_3). Each decision tree represents the discrete set of feasible decisions available to the lower-level decisionmaker(s), whereas the utilities represented at the leaf nodes of the tree respectively depend on the upper-level persuader's actions via the indicated functions.

The persuader affects the decision problem prior to prospect selection and desires the decisionmaker to choose prospect B subject to the constraints listed. For an interested reader, the Appendix details an explicit model of this influence setting. Although this particular instance can be solved by inspection (e.g., $a_1 = a_2 = 0$, $a_3 > 216.67$ is one of many alternative optimal solutions), it provides a simple illustration of influence in a decision analysis setting. Likewise, such influence problems can be represented as bilevel programs:

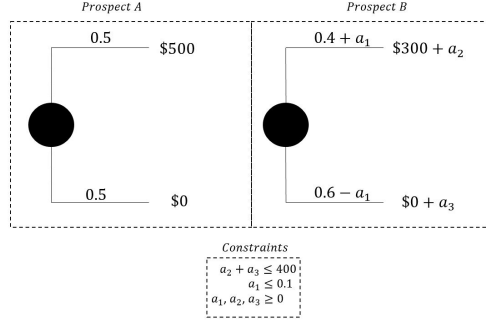


Figure 2. Simple Influence Example under Risk for Risk-Neutral, EUT Decisionmaker

$$\begin{aligned}
 & \max_{\mathbf{a}} \Phi \\
 & \text{subject to } \mathbf{a} \in A, \\
 & \max_{\zeta, \Phi} \sum_{j \in J} \zeta_j EU(j), \\
 & \text{s.t. Constraints (1a) – (1b),} \\
 & EU(j) = f_1(\mathbf{a}), \forall j \in J, \\
 & \Phi = f_2(\zeta), \\
 & \Phi \in \{0, 1\}.
 \end{aligned}$$

Within the bilevel programming framework, the lower player is the individual being persuaded and who selects the prospect maximizing their expected utility. The upper-level player is able to affect the individual's expected utility through some action \mathbf{a} from a larger action space A to maximize the binary persuasion indicator variable Φ . At optimality, the prospect corresponding to $\zeta_j^* = 1$ is the prospect selected by the decisionmaker, and a value of $\Phi^* = 1$ indicates that such a prospect is the one preferred by the persuader.

The persuader's actions are designed to induce some selection, indicating the modeling of choice must be governed by the decisionmaker's psychology. However,

it is well documented that without assistance, humans violate the Expected Utility paradigm (e.g., see Andreoni and Sprenger, 2010). As evidenced by the constant consequence and constant ratio paradoxes, decisionmakers do not always make choices that maximize expected utility. Therefore, the corresponding behavioral decision problem of an individual choosing between $|J|$ options can be represented as

$$\max_{\zeta} \sum_{j \in J} \zeta_j DU(j)$$

subject to Constraints (1a) – (1b),

wherein $DU(j)$ represents the subjective utility of a prospect $j \in J$ under some descriptive theory of choice (e.g., Kontek and Lewandowski, 2017)

Given the bounded rationality exhibited in empirical psychological testing, we contend a descriptive theory of choice is most appropriate to model human persuasion. Utilizing such a descriptive model, the problem of influencing an individual to adopt a desired course of action can be represented as follows:

$$\max_{\mathbf{a}} \Phi$$

subject to $\mathbf{a} \in A$,

$$\max_{\zeta, \Phi} \sum_{j \in J} \zeta_j DU(j),$$

s.t. Constraints (1a) – (1b),

$$DU(j) = g_1(\mathbf{a}), \forall j \in J,$$

$$\Phi = g_2(\zeta),$$

$$\Phi \in \{0, 1\}.$$

For the remainder of this research, Cumulative Prospect Theory (CPT) is utilized to model the descriptive utility, $DU(j)$, for reasons described in Section 2.2 (Tversky

and Kahneman, 1992). Under this theory, a prospect f having n gain outcomes and m loss outcomes, indexed separately as 1 through n and $-m$ through -1 , each with value x_k from the reference point is evaluated via CPT as

$$\begin{aligned}
 V(f) &= V^+(f) + V^-(f), \text{ where} \\
 V^+(f) &= \sum_{k=1}^n \pi_k^+ v(x_k), \\
 V^-(f) &= \sum_{k=-m}^0 \pi_k^- v(x_k), \\
 \pi_n^+ &= W^+(A_n), \\
 \pi_{-m}^- &= W^+(A_{-m}), \\
 \pi_k^+ &= W^+(A_k \cup \dots \cup A_n) - W^+(A_{k+1} \cup \dots \cup A_n), \quad 0 \leq k \leq n-1 \\
 \pi_k^- &= W^+(A_{-m} \cup \dots \cup A_k) - W^+(A_{-m} \cup \dots \cup A_{k-1}), \quad 1-m \leq k \leq 0.
 \end{aligned}$$

All outcomes A_k are indexed in ascending order of their corresponding x_k -value. W^+ and W^- are the event weighting functions for gains and losses, respectively. The values π^+ and π^- are the decision weights utilized to determine the respective component gains, $V^+(f)$, and component losses, $V^-(f)$, in conjunction with a utility function $v(\cdot)$ that is concave for gains and convex for losses.

Whereas CPT is often viewed as a theory of decision making under risk, Tversky and Fox (1995) illustrated its applicability to decision making under uncertainty, and that the uncertainty weighting function demonstrated similar behavior to the probability weighting function. Moreover, Fox and Tversky (1998) demonstrated that decisions under uncertainty can be modeled suitably using a two-stage method. In the first stage, probability judgments comporting with support theory (Tversky and Koehler, 1994) are first ascertained to characterize and quantify the uncertainty and, in the second stage, these results are substituted into the original CPT framework that accounts for risk. Utilizing these results, we incorporate CPT into our mathe-

mathematical programming representations of persuasion under both conditions of risk and uncertainty.

Our modeling framework unifies and augments three separate strands of research: decision analysis problems, cheap talk games, and multi-level programs. Within the context of the Decision Analysis discipline, our work considers a self-interested advisor (i.e., persuader) as opposed to one who is altruistic. It is unique in the realm of cheap talk games to allow for multiple, lower-level players who assess the value of an outcome via Cumulative Prospect Theory; to address uncertain outcomes; and to enable the upper-level players to adopt a wider range of actions. Finally, the modeling framework is original in the context of multi-level programming by modeling a principal-agent(s) problem with the incorporation of behaviorist methods applied to the lower-level players in a stochastic, Stackelberg framework.

Therefore, in a manner akin to Fry and Binner (2016) and Becker-Peth and Thonemann (2016), our research also furthers the Behavioral Operations Research paradigm described in Hämäläinen et al. (2013) with respect to incorporating agent psychology in problem solving situations. More specifically, this work advances the Behavioral Operations Research thread referenced by Franco and Hämäläinen (2016) that “concentrates on the use of the OR approach to model human behavior in complex settings” by combining quantitative psychology theories with mathematical programming. In this way, our work is akin to that of Shi and Lian (2016), Keller and Katsikopoulos (2016), and Argyris and French (2017) who adopt a behavioral perspective in queuing theory, heuristic optimization, and multi-criteria decision analysis applications, respectively, as our research attempts to address the task outlined in Becker (2016) of developing a “close connection with the core disciplines of OR”.

The remaining of this paper is structured as follows. In Section 2.2, we provide an overview of scholarly research pertaining to persuasion and also explain why

CPT is selected as the preferred theory of descriptive utility. In Section 2.3, we develop single-level mathematical programming formulations of the descriptive bilevel persuasion problem corresponding to situations exhibiting risk and uncertainty, respectively. Therein, we establish the soundness of our modeling methodology, as supported by literature that quantitatively documents experimental evidence of the underlying theories, and which portends the ability to parameterize instances of our models. In Section 2.4, we provide illustrative instances of the models' use: one under risk, one under uncertainty with an evidence strength metric, and one under uncertainty without an evidence strength metric. Section 2.5 discusses the implications of this research, its limitations, and the potential for application.

2.2 Relevant Literature

We begin by providing a multi-disciplinary review of persuasion literature to illustrate the relevance of our models and emphasize their complementary nature to previous studies. The section concludes with a discussion of alternative frameworks for decision making under uncertainty and defends subsequent modeling choices.

Persuasion.

Persuasion has been widely studied in the field of social psychology. Gass and Seiter (2015) defined persuasion as an “activity of creating, reinforcing, modifying or extinguishing beliefs, attitudes, intentions, motivations, and/or behaviors within the constraints of a given communication context”. Thus, persuasion is not limited to speech, as it may include any number of action types. Scholars in this field have extensively studied persuasion and propaganda. Research conducted immediately following World War I attempted to describe the effects of messaging on attitude (Jowett and O'Donnell, 2014). During World War II emphasis was placed on understanding

propaganda, counterpropaganda, attitudes, and persuasion (Jowett and O'Donnell, 2014).

In recent history, the focus shifted alternatively to predicting or altering future behavior (Jowett and O'Donnell, 2014). Many formal theories of persuasion have been presented. Two prevailing models in social psychology are the Elaboration Likelihood Method and the Heuristic Systematic Method (Gass and Seiter, 2015). The Elaboration Likelihood Method postulates two distinct avenues of persuasion: central and peripheral processing. Central processing encapsulates conscious deliberation of a message whereas peripheral processing involves the effect of stimuli not directly related to a message (e.g., a salesperson's pitch vis-a-vis their physical appearance). The Heuristic Systematic Method also proposes that persuasion occurs through two avenues (i.e., deliberate systematic processing and instinctive heuristic processing) but contends that messages can traverse both paths simultaneously.

Related ideas in social psychology are the Theory of Cognitive Dissonance, Exposure Theory, and the Theory of Planned Behavior. The Theory of Cognitive Dissonance postulates that humans desire a state of cognitive consistency in thought and will change opinions and/or behavior in order to achieve it (Festinger, 1957). In contrast, Exposure Theory propounds that people who are highly exposed to an idea are more likely to accept it. The content or quality of the idea is not nearly as important as the frequency of exposure; the mere comfort of familiarity is the convincing factor (Zajonc, 2001). Finally, the Theory of Planned Behavior hypothesizes that a person's choices of behavior are influenced by societal pressures; individual perception, power, and attitude; and other factors facilitating or inhibiting the choice (Ajzen, 1991).

Scholars in the field of Communications focus less on the mental processing of persuasion and more on the features of the messaging itself. Shen and Bigsby (2013) presented a thorough literature review regarding the effects of message content, struc-

ture, and style on persuasion and provided an overview of how these features interact. The effect of messaging is also well studied in Economics, particularly with regard to games of incomplete information. It is common in such games for players to engage in signaling, an act of providing or receiving information which may be costly or, in *cheap talk*, a similar but cost-free information sharing action. In this setting, emotions of the players are generally disregarded, and perspectives are manipulated according to the tenants of Expected Utility Theory. Signaling games and cheap talk are covered extensively by Shoham and Leyton-Brown (2008) with a unique focus on natural language processing. Chakraborty and Harbaugh (2010) and Croson et al. (2003) discussed applications of cheap talk and its potential consequences in persuasion and bargaining games.

With a view towards applying persuasion operations, agents often attempt to manipulate subjective and emotional assessments to gain support. For instance, recruiting efforts by terrorist organizations usually involve a radicalization process wherein recruits' political or world views are shaped to align with the group's interests. McCauley and Moskalenko (2008) described a variety of tactics utilized for radicalization such as isolation, personal victimization, and the use of political grievances. Other behavioral modification campaigns are central to Military Information Support Operations (MISO) designed to influence the behavior of foreign governments, organizations, groups, or individuals (Boyd, 2011). To implement a successful persuasion operation, an agent must understand the effects of one's actions on a decisionmaker. However, such knowledge is not only useful in the development of offensive messaging, but it also provides insight regarding how to counter persuasion operations (e.g., counter-terrorist organizations, influence resistance campaigns).

The field of Marketing has a long history of studying how to sway opinions. Belk (2007) described the evolution of commercial marketing research from a focus on de-

scriptive data to qualitative research. The qualitative research the author describes is focused on understanding the underlying motivation for consumer behavior. However, the qualitative study of persuasion is ill-equipped to solve strategic marketing allocation problems. In this vein, Cetin and Esen (2006) tailored the weapon target assignment model to marketing campaigns, but they focused on budget and scheduling constraints while largely abstaining from psychological factors. Therefore, the next evolution in commercial marketing research is arguably a transition from qualitative to quantitative influence modeling. Our research provides the initial models to make such an advance with the necessary incorporation of qualitative psychological effects. Moreover, our research provides the impetus to perform the psychological and physiological experiments necessary to make this advance (e.g., see the consumer research described in Ares and Varela, 2018).

Similar concepts to those discussed by Belk (2007) have been applied to Social Marketing, wherein the underlying goal is to influence behavioral patterns within a society for the greater good. Kotler and Zaltman (1971) described the power and limitations of such applications and explained how social marketing has been conducted by churches to increase their attendance, by charities to generate income, and even by symphonies to draw a larger audience. Many non-profit groups have tried to engineer such social change against drug or tobacco use in the United States (Truth Initiative, 2018). Research in criminology has advocated the use of similar social marketing tactics to reduce crime of varying types (Homel and Carroll, 2009). With regard to politics, the role of emotion in campaigns has been studied by Brader (2005), wherein he determined that emotional appeals in campaign ads can affect voter behavior. Brader (2005) concluded that “emotionally evocative ads do not simply sway voters directly, but change the manner in which voters make choices”. These overarching changes in decision making are part of what we seek to model within a prescriptive,

quantitative approach to the problem of persuasion.

Other Models for Decision Making Under Uncertainty.

We choose to utilize CPT as our framework for decision-making under conditions of either risk or uncertainty. This choice is driven by (1) the interpretation of CPT as a characteristic of intuitive thinking (Kahneman, 2011) and the deep heuristic-and-bias decision-making literature describing how to affect such thought processes (Kahneman and Tversky, 1972, 1981; Park and Lessig, 1981; Tversky and Kahneman, 1981; Taylor, 1982), and (2) CPT's applicability to risk and uncertainty (Kahneman and Tversky, 1979; Fox and Tversky, 1998).

Other models that challenge the preeminence of CPT as a descriptive model were considered but set aside. For example, the Adaptive Toolbox proposed by Gigerenzer and Selten (2002) purports that a collection of situational specific cognitive heuristics are utilized when acting under uncertainty. Historically, the Adaptive Toolbox framework has been the leading competitor of CPT. However, we choose to not utilize this framework due to the disaggregated, environment-specific nature of its models. CPT provides a single mathematical model that can be applied without regard to the underlying decision environment, whereas the Adaptive Toolbox provides a variety of methods to be utilized in differing environments. Furthermore, Pachur et al. (2017) demonstrated how a variety of boundedly rational heuristics exhibit characteristic CPT parameter profiles. If a decisionmaker is assumed to use such a heuristic decisionmaking process, their behavior can be incorporated into our models by utilizing the appropriate profile.

Other, less established alternatives to CPT exist, but they are not well suited for a general persuasion model. The Decision-by-Sampling (DbS) model proposed by Stewart et al. (2006) describes decision making as a process of binary, ordinal com-

parisons of an attribute with past experiences; it does not assume a stable probability weighting or utility function, and argues its comparative framework is the foundation of CPT-characteristic behavior. DbS supports the notion of affecting choice via persuasion, but it may restrict actions to only those based on the Exposure Theory of messaging. More recently, Kontek and Lewandowski (2017) proposed Range Dependent Utility (RngDU) as a model describing the same decision making paradoxes addressed by CPT, albeit without the use of probability weighting. Instead, each lottery in a specified range of low and high outcomes possesses its own utility function, and a single decision utility function can be derived by linearly scaling the utility of all lotteries (i.e., with the lowest possible utility having a value of zero and the highest utility having a value of one). The linearity-in-probability is a convenient feature for calculation, and the authors claim RngDU to be a better predictive model for multi-outcome prospects compared to CPT by providing comparative examples from the literature. However, the RngDU theory is relatively nascent and does not yet incorporate loss aversion. In contrast, the mathematical framework presented in CPT provides a well-established and generally accepted quantitative manner to describe psychological and qualitative phenomena related to persuasion that can be adapted to accommodate other descriptive theories of choice.

2.3 Influence Modeling

Herein, we set forth the mathematical programming formulation describing how a persuader can optimally interact with a subset of individuals based upon assumptions of how persuasive actions affect each prospect, as well as respective individuals' evaluation of the prospects.

To provide proper context prior to introducing the model, we detail the literature supporting the viability of quantifying the parameter update functions included

within our modeling framework (e.g., general functional forms to represent the effect of persuasion on prospect characteristics or decisionmaker risk attitude) which, in turn, establishes that such update functions can be realized. Often, the components comprising an individual's risk attitude (e.g., the probability weighting and utility functions) are assumed to be static, cognitive features. However, modern psychological and neuroeconomic studies have shown risk attitudes to be dynamic traits affected by emotional, pharmacological, and other context-dependent factors (e.g., Kugler et al., 2012; Lempert and Phelps, 2014; Stewart et al., 2015). The following section introduces readers to related parametric-oriented research before providing the General Persuasion Program (GPP) variants for conditions of risk and uncertainty.

The Parameter Update Functions.

Two fundamental assumptions underlying our GPP models are (1) a persuader's actions can influence an individual's perception of the probability and utility of outcomes for a given prospect, and (2) such impacts can be quantified. In this section, we address the literature that supports the quantification of such effects which, in turn, establishes that these effects can be realized. Given a thorough review of findings in the literature that are supported by well-conducted human subjects research, we find evidence that influence on the perception of an outcome's likelihood, the perception of an outcome's value, and an individual's cumulative prospect theory parameters each can be quantified.

Of course, if a persuader has direct control over the source of risk (uncertainty) or the event outcomes (e.g., a casino), the update functions for these parameters are readily apparent. However, a persuader may still be able to influence how individuals perceive an event's likelihood or its utility in situations wherein there exists no formal, direct control mechanism. We find justification in the psychological literature

pertaining to human evaluation of uncertainty and decision-making. Lerner et al. (2015) detailed extensively the psychological literature that empirically shows the effects of emotion on risk perception. For example, fearful people tend to see greater risk whereas angry people perceive less risk, and prideful individuals view gains as in their control whereas surprised individuals view gains as unpredictable (Lerner and Keltner, 2000; Lerner et al., 2015). Furthermore, in the heuristic-and-bias approach of statistical reasoning (Kahneman and Tversky, 1979; Kahneman, 2011), the affect, representativeness, and availability heuristics, as well as their derivatives (e.g., the simulation and fluency heuristics), demonstrate how superfluous information can alter the perceived likelihood of an event or the value of its outcome (Kahneman and Tversky, 1981; Hertwig et al., 2008; Kahneman, 2016). In a persuasion setting, the source of this information is simply a self-interested external agent. The *Tom W.* experiments of Kahneman and Tversky (1973) demonstrate such an interaction through representativeness; the authors affect the participants' perception of probability by controlling the stimuli used to make this judgment.

With specific regard to the probability weighting and utility functions, Stewart et al. (2015) observed statistical evidence concerning their malleable shapes. The authors found that the “manipulation of the distribution of gains [and] risks... systematically changes the utility, weighting, and discounting functions”, and they concluded that these functions are not stable but are instead context dependent. As the probability weighting and utility functions collectively describe risk attitude, their results also imply its malleability. This is confirmed by Kugler et al. (2012) who demonstrated that fear and anger alter an individual's risk attitude, by Ariely and Loewenstein (2006) who found a similar effect in men due to sexual arousal, and by numerous economists who have examined the effect of past losses on the risk attitudes of stock brokers (e.g., Thaler, 1985; Weber and Zuchel, 2005; Liu et al., 2010). How-

ever, with few exceptions, such changes are not quantified in terms of the associated CPT-parameters (Campos-Vazquez and Cuijly, 2014; Schulreich et al., 2014)

When affecting decision under uncertainty, the belief-based model by Fox (1999) illustrates that the probability weighting function parameters may be affected through a source preference appeal. The authors provided evidence of the competence hypothesis such that ambiguity aversion is reversed in situations wherein a person feels knowledgeable of the source. It is suggested that this effect can be accounted for by varying a parameter in the probability weighting function. Therefore, if the persuader has some control over the source of uncertainty, they may be able to affect the weighting function parameters by appealing to an individual's preference. Similar results could be accomplished by adapting an individual's perception of the source of uncertainty, perhaps through education.

Likewise, since the "loss aversion [parameter] is volatile and depends much on framing" (Wakker, 2010), there exists empirical evidence of its malleability. Plott and Zeiler (2005) were able to make it disappear completely when participants were explicitly instructed on loss aversion and framing effects. Gächter et al. (2007) showed that loss aversion increases with age, wealth, and income but decreases with education. Furthermore, the correlation found with respect to activity in areas of the brain when confronted with gains and losses by Tom et al. (2007) suggest that, if these areas can be stimulated when a participant is facing a decision, perhaps loss aversion can be altered. Moreover, Lerner et al. (2004) found emotional stimuli to reverse the endowment effect, implying emotions affect either the loss aversion parameter or the frequency of outcome evaluation.

With regard to reference points, their manipulation depends largely on the decision context. For example, if the decision to be made is measured in monetary units, the persuader may be able to move the reference point by instituting fines or payments

before a lottery is selected. In general, the perception of reference points is highly susceptible to framing effects which can be controlled, to some extent, by a persuader (Tversky and Kahneman, 1981; Wakker, 2010).

Moreover, if an Adaptive Toolbox heuristic (e.g., Gigerenzer and Selten, 2002) is assumed to be employed for a decision problem, the results of Pachur et al. (2017) imply the malleability of CPT parameters. That is, if a decisionmaker can be compelled to switch from one heuristic to another, the parameter update functions can map the change from one heuristic's CPT parameter profile to another.

Ultimately, risk and value perception, as well as an individual's CPT parameters, are mathematical constructs designed to capture cognition. From a neuroeconomic perspective, they all arise from neuronal interactions in a dynamic brain learning from a dynamic world. Therefore, the hypothesis that these elements are static and not malleable is incongruent with the field's underlying tenets. Such a perspective is documented with numerous works published within *Neuroeconomics: Decision Making and the Brain* edited by Paul Glimcher and Ernest Fehr, to include studies focusing on the effects of emotions and pharmacology in decisions under risk and uncertainty (Lempert and Phelps, 2014; Crockett and Fehr, 2014).

Acknowledging that such a mapping is feasible, as evidenced by the quantifiable impacts of persuasion on decisionmakers' CPT parameters via the published literature, we seek to demonstrate the utility of such mappings via our mathematical programming formulations and several use cases in Section 2.4.

Proposed Model under Risk.

Herein, we develop a single level representation of the aforementioned bilevel math programming formulation to represent the effects of persuasion on a human decision-maker under conditions of risk with a generic upper level objective function. This

model is general in that it concerns influencing multiple decisionmakers, each of whom may have different values for their Cumulative Prospect Theory parameters, to prefer a certain prospect (i.e., become supporters). In this setting, a persuader may have multiple, competing objectives; it may be desired to achieve a threshold of support at minimum cost, to maximize the number of supporters, to optimize a weighted combination of both objectives, or any number of variants. However, the majority of these objectives are concerned with the number of supporters and the persuasion actions utilized. As such, we represent the objective function within our mathematical programming formulation as a generic function of these variables. Likewise, our general model allows for the malleability of Cumulative Prospect Theory parameters via persuasion operations, and so the perceived outcomes and probabilities of individual prospects can be affected accordingly. Our models, therefore, can accommodate a range of persuasive actions, to include the intentional triggering of the *incidental affect* wherein an emotion unrelated to the decision can shift choices. This phenomena is well document in the neuroscience literature but, to our knowledge, has not been adapted to quantitative decision models (Lempert and Phelps, 2014). The models presented are therefore versatile; however, if one disagrees with the malleability of a parameter subset, it can be assumed constant and the optimization formulations are still applicable. Such applications are illustrated in Section 2.4.

We begin by introducing the requisite sets, parameters, decision variables, and intermediary decision variables. In defining the model, we use the single parameter weighting function and utility functions originally set forth by Tversky and Kahneman (1992). However, other probability weighting functions or utility functions can readily be used by substituting their functional forms and parameters in the corresponding constraints.

Sets

I : Set of targeted members upon which persuasion is being conducted of the form $I = \{1, 2, \dots\}$

J_i : Set of n_i prospects offered to each member i of the form $J_i = \{1, \dots, n_i\}$

K_{ij} : Set of m_{ij} events (i.e., outcomes) for member i in prospect j of the form $K_{ij} = \{1, \dots, m_{ij}\}$

A : Space of feasible persuasion actions

Parameters

q_i : Preferred prospect in J_i that persuader wants member i to support

\hat{y}_{ijk} : Baseline raw value for k^{th} event of prospect j for member i before persuasion

\hat{p}_{ijk} : Baseline probability of k^{th} event of prospect j for member i before persuasion

$\hat{\gamma}_i$: Baseline gain probability weighting coefficient for member i before persuasion

$\hat{\delta}_i$: Baseline loss probability weighting coefficient for member i before persuasion

$\hat{\alpha}_i$: Baseline gain utility coefficient for member i before persuasion

$\hat{\beta}_i$: Baseline loss utility coefficient for member i before persuasion

$\hat{\lambda}_i$: Baseline loss aversion coefficient for member i before persuasion

\hat{r}_i : Baseline reference point for member i before persuasion

Δ : Discriminable Factor for favored prospect

M : Arbitrary, sufficiently large real number

We assume a set of members, I , are offered a set of prospects, J_i . Members of I can be individuals or aggregate demographic populations, depending on the setting. The set J_i has n_i prospects which may or may not be unique to each member. Each prospect $j \in J_i$ has m_{ij} potential events. The effect on the editing phase on these events is determined *a priori*. Also, the value of an event before persuasion as discerned from some measure is represented by \hat{y}_{ijk} and its associated probability by \hat{p}_{ijk} .

The persuader desires for each member i to select a prospect q_i . By taking ac-

tion(s), represented by the vector $\mathbf{a} \in A$, the persuader is able to potentially affect the outcomes and probabilities for each prospect, or alter a member's baseline reference point, loss aversion, probability weighting, and utility curvature parameters, depending on the specific assumptions of the functional forms that adapt our general model to a specific context.

Primary Decision Variables

\mathbf{a} : Vector of influence operations conducted

$T_{ijkk'}^+$: Equal to 1 if the k^{th} event of prospect j faced by member i is the $(m_{ij} - 1 + k')^{th}$ greatest gain among events in K_{ij} , and 0 otherwise, defined for all (i, j) combinations

$T_{ijkk'}^-$: Equal to 1 if the k^{th} event of prospect j faced by member i is the $(k')^{th}$ greatest loss among events in K_{ij} , and 0 otherwise, defined for all (i, j) combinations

Φ_i : Binary indicator equal to 1 if member i strongly favors prospect q_i , 0 otherwise

$z_{ij}^{q_i}$: Binary indicator for $q_i \neq j$ that equals 1 if member i prefers prospect q_i to prospect j , 0 otherwise

s_{ij}^{pos} : Positive portion of difference in prospect values q_i and j for member i

s_{ij}^{neg} : Negative portion of difference in prospect values q_i and j for member i

Ψ_{ij} : Binary variable enforcing that both s_{ij}^{pos} and s_{ij}^{neg} cannot be positively-valued

CPT relies upon a rank-based ordering of each prospect's outcomes to properly calculate its value. If we allow the persuader's actions to change outcome values, their probabilities, and a member's Cumulative Prospect Theory parameter values, then the original ordering of a prospect's events can change. As such, a mapping is developed to ensure the events are indexed in ascending order of their outcome values for use in the probability weighting functions. The binary variables $T_{ijkk'}^+$ form a matrix mapping of unordered gains to ordered gains, and $T_{ijkk'}^-$ a mapping of unordered losses to ordered losses. Applied simultaneously, $T_{ijkk'}^+$ and $T_{ijkk'}^-$ create a

mapping that sorts the updated gains and losses in ascending order.

The variable Φ_i equals one if a member i strictly prefers prospect q_i to all others. A strict preference is induced versus a weak preference (i.e., strict versus weak inequality) to ensure that members who are indifferent among q_i and some other subset of prospects are not misrepresented as supporters via an optimistic assumption by the persuader. Members demonstrating such indifference may or may not be supportive, and we assume a conservative persuader will presume them to be the latter. Likewise, the discriminable factor Δ enables the persuader to specify a minimum difference value between prospects such that a member is identified as a supporter. Binary comparisons between prospects are represented via $z_{ij}^{q_i}$; this binary variable equals one if member i prefers prospect q_i to prospect j , and zero otherwise. The variables s_{ij}^{pos} and s_{ij}^{neg} are utilized to quantify either the positive or negative difference between prospect q_i and prospect j , wherein $s_{ij}^{pos} > 0$ indicates prospect q_i is preferred and $s_{ij}^{neg} > 0$ indicates the converse is true. Finally, Ψ_{ij} is utilized in the ensuing math programming formulation to ensure that only s_{ij}^{pos} or s_{ij}^{neg} may be positively valued.

Intermediate Decision Variables

$f_{ijk}(\mathbf{a})$: Persuasion effect on event k^{th} raw outcome value for prospect j and member i

$g_{ijk}(\mathbf{a})$: Persuasion effect on event k^{th} probability for prospect j and member i

$h_i^\theta(\mathbf{a})$: Persuasion effect on curvature, distortion, loss aversion or reference point parameters for member i

x_{ijk} : Gain/loss for member i for k^{th} event of prospect j after persuasion

p_{ijk} : Probability of k^{th} event of prospect j for member i after persuasion

γ_i : Gain probability weighting coefficient for member i after persuasion

δ_i : Loss probability weighting coefficient for member i after persuasion

α_i : Gain utility curvature coefficient for member i after persuasion

β_i : Loss utility curvature coefficient for member i after persuasion

λ_i : Loss aversion coefficient for member i after persuasion

r_i : Reference point for member i after persuasion

$t_{ijk'}^+$: Ascending rank based list of x_{ijk} gains corresponding with mapping $T_{ijkk'}^+$

$t_{ijk'}^-$: Ascending rank based list of x_{ijk} losses corresponding with mapping $T_{ijkk'}^-$

$b_{ijk'}^+$: Corresponding probabilities for sorted $t_{ijk'}^+$ events

$b_{ijk'}^-$: Corresponding probabilities for sorted $t_{ijk'}^-$ events

$\pi_{ijk'}^+$: Gain decision weight for member i for k' th event of prospect j after persuasion

$\pi_{ijk'}^-$: Loss decision weight for member i for k' th event of prospect j after persuasion

The intermediate decision variables are functions of the primary decision variables. The variables x_{ijk} and p_{ijk} respectively represent the updated events, in terms of gains or losses from the updated reference point, and the updated probabilities of these events after persuasion. Likewise, the variables γ_i , δ_i , α_i , β_i , λ_i , and r_i are the probability weighting, utility curvature, loss aversion, and reference point values after persuasion, respectively. The effects of persuasion on outcomes, probabilities, and prospect theoretic parameters are correspondingly represented by the values f_{ijk} , g_{ijk} , and h_i^θ , where θ is an index on the functional parameters γ_i , δ_i , α_i , β_i , λ_i , and r_i . The list of updated gains and losses in ascending order is represented by $t_{ijk'}^+$ and $t_{ijk'}^-$, respectively, with associated probabilities $b_{ijk'}^+$ and $b_{ijk'}^-$. The last six intermediate decision variables are drawn from Tversky and Kahneman (1992), wherein $\pi_{ijk'}^+$ and $\pi_{ijk'}^-$ are the decision weight values derived from the probability weights $w^+(b_{ijk'}^+, \gamma_i)$ and $w^-(b_{ijk'}^-, \delta_i)$, and $V^+(x_{ijk'}, \alpha_i)$ and $V^-(x_{ijk'}, \beta_i)$ are the values of updated gains and losses, respectively.

As previously mentioned, other probability weighting functions can be utilized if their parameters and functional forms are substituted as appropriate. However,

the decision regarding which probability weighting function is most appropriate depends on the modeling context. If members are aggregate populations, the Tversky and Kahneman (1992) form can be expected to estimate the weighting function well (Gonzalez and Wu, 1999). However, under the assumption that members of the set I are individuals instead of groups, a two-parameter probability weighting function such as the Linear-in-Log-Odds or Prelec's compound invariance functions may be most appropriate (Cavagnaro et al., 2013).

General Persuasion Program (GPP) Formulation

$$\max f(\Phi_i, \mathbf{a}) \quad (2a)$$

subject to

Parameter Update Constraint Set

$$x_{ijk} + r_i = \hat{y}_{ijk} + f_{ijk}(\mathbf{a}), \quad \forall i \in I, j \in J_i, k \in K_{ij}, \quad (2b)$$

$$p_{ijk} = \hat{p}_{ijk} + g_{ijk}(\mathbf{a}), \quad \forall i \in I, j \in J_i, k \in K_{ij}, \quad (2c)$$

$$\theta_i = \hat{\theta}_i + h_i^\theta(\mathbf{a}), \quad \forall i \in I, \theta = \{\gamma, \delta, \alpha, \beta, \lambda, r\}, \quad (2d)$$

Ordering Constraint Set

$$\sum_{k=1}^{m_{ij}} T_{ijkk'}^+ + T_{ijkk'}^- = 1, \quad \forall i \in I, j \in J_i, \quad (2e)$$

$$\sum_{k'=1}^{m_{ij}} T_{ijkk'}^+ + T_{ijkk'}^- = 1, \quad \forall i \in I, j \in J_i, \quad (2f)$$

$$t_{ijk'}^+ = \sum_{k=1}^{m_{ij}} T_{ijkk'}^+ x_{ijk}, \quad \forall i \in I, j \in J_i, k' \in K_{ij}, \quad (2g)$$

$$t_{ijk'}^- = \sum_{k=1}^{m_{ij}} T_{ijkk'}^- x_{ijk}, \quad \forall i \in I, j \in J_i, k' \in K_{ij}, \quad (2h)$$

$$t_{ij(k'+1)}^- \geq t_{ijk'}^-, \quad \forall i \in I, j \in J_i, k' \neq m_{ij}, \quad (2i)$$

$$t_{ij(k'+1)}^+ \geq t_{ijk'}^+, \quad \forall i \in I, j \in J_i, k' \neq m_{ij}, \quad (2j)$$

$$t_{ijk'}^+ \geq 0, \quad \forall i \in I, j \in J_i, k' \in K_{ij}, \quad (2k)$$

$$t_{ijk'}^- \leq 0, \quad \forall i \in I, j \in J_i, k' \in K'_{ij}, \quad (2l)$$

$$b_{ijk'}^+ = \sum_{k=1}^{m_{ij}} T_{ijkk'}^+ P_{ijk}, \quad \forall i \in I, j \in J_i, k' \in K_{ij}, \quad (2m)$$

$$b_{ijk'}^- = \sum_{k=1}^{m_{ij}} T_{ijkk'}^- P_{ijk}, \quad \forall i \in I, j \in J_i, k' \in K_{ij}, \quad (2n)$$

CPT-Value Constraint Set

$$\pi_{ijm_{ij}}^+ = w^+(b_{ijm_{ij}}^+, \gamma_i), \quad \forall i \in I, j \in J_i, \quad (2o)$$

$$\pi_{ijk'}^+ = w^+\left(\sum_{l=k'}^{m_{ij}} b_{ijl}^+, \gamma_i\right) - w^+\left(\sum_{l=k'+1}^{m_{ij}} b_{ijl}^+, \gamma_i\right), \quad \forall i \in I, j \in J_i, k \neq m_{ij}, \quad (2p)$$

$$\pi_{ij,1}^- = w^-(b_{ij1}^-, \delta_i), \quad \forall i \in I, j \in J_i, \quad (2q)$$

$$\pi_{ij,k'}^- = w^-\left(\sum_{l=1}^{k'} b_{ijl}^-, \delta_i\right) - w^-\left(\sum_{l=1}^{k'-1} b_{ijl}^-, \delta_i\right), \quad \forall i \in I, j \in J_i, k' \neq 1, \quad (2r)$$

$$\begin{aligned} & \left(\sum_{k'=1}^{m_{ij}} \pi_{iq_i k'}^+ V^+(t_{iq_i k'}^+, \alpha_i) + \pi_{iq_i k'}^- V^-(t_{iq_i k'}^-, \beta_i, \lambda_i) \right) \\ & - \left(\sum_{k'=1}^{m_{ij}} \pi_{ij k'}^+ V^+(t_{ij k'}^+, \alpha_i) + \pi_{ij k'}^- V^-(t_{ij k'}^-, \beta_i, \lambda_i) \right) - s_{ij}^{pos} + s_{ij}^{neg} - \Delta = 0, \end{aligned} \quad (2s)$$

$$\forall i \in I, j \in J_i,$$

Preferred Prospect Constraint Set

$$s_{ij}^{pos} \leq M(1 - \Psi_{ij}), \quad \forall i \in I, j \in J_i, \quad (2t)$$

$$s_{ij}^{neg} \leq M\Psi_{ij}, \quad \forall i \in I, j \in J_i, \quad (2u)$$

$$Ms_{ij}^{pos} \geq z_{ij}^{q_i}, \quad \forall i \in I, j \in J_i, \quad (2v)$$

$$(n_i - 1)\Phi_i \leq \sum_{j \neq q_i} z_{ij}^{q_i}, \quad \forall i \in I, \quad (2w)$$

Decision Variable Domain Constraint Set

$$\mathbf{a} \in A, \quad (2x)$$

$$\Phi_i \in \{0, 1\}, \quad \forall i \in I, \quad (2y)$$

$$z_{ij}^{q_i}, \Psi_{ij} \in \{0, 1\}, \quad \forall i \in I, j \in J_i, \quad (2z)$$

$$s_{ij}^{pos}, s_{ij}^{neg} \geq 0, \quad \forall i \in I, j \in J_i, \quad (2aa)$$

$$T_{ijkk'}^+, T_{ijkk'}^- \in \{0, 1\}, \quad \forall i \in I, j \in J_i, k \in K_{ij}, k' \in K_{ij}, \quad (2ab)$$

where $\forall i \in I, j \in J_i, k' \in K_{ij}$ we have

$$w^+(b_{ijk'}^+, \gamma) = \frac{(b_{ijk'}^+)^\gamma}{((b_{ijk'}^+)^\gamma + (1 - b_{ijk'}^+)^\gamma)^{\gamma-1}}, \quad (3a)$$

$$w^-(b_{ijk'}^-, \delta) = \frac{(b_{ijk'}^-)^\delta}{((b_{ijk'}^-)^\delta + (1 - b_{ijk'}^-)^\delta)^{\delta-1}}, \quad (3b)$$

$$V^+(t_{ijk'}^+, \alpha_i) = (t_{ijk'}^+)^{\alpha_i}, \quad (3c)$$

$$V^-(t_{ijk'}^-, \beta_i, \lambda_i) = -\lambda(-t_{ijk'}^-)^{\beta_i}. \quad (3d)$$

The objective function (2a) maximizes a function of the number of supporters and the persuasion actions taken, where the specific form of the function depends upon the context of the problem instance. The constraints can be grouped as follows: the Parameter Update (Constraints (2b)–(2d)), Ordering (Constraints (2e)–(2n)), CPT-Value (Constraints (2o)–(2s)), Preferred Prospect (Constraints (2t)–(2w)), and Decision Variable Domain (Constraints (2x)–(2ab)) Constraint Sets. More formally, Constraints (2b) and (2c) ensure that, for each member, for every event in their set of prospects, the event’s value and associated probability are updated in accordance with the persuasion action. Constraint (2d) updates each member’s probability distortion, value curvature, loss aversion, and reference point values with respect to the same persuasion action. This research assumes the independence of parameter updates and that they are functions of only the persuasion action. The exact natures of these functions are not known, as the possibility of such updates has only recently been studied in the literature. Moreover, such parameter updates may be interrelated. For example, the function to update the loss probability distortion parameter may depend on the action and the previous gain probability distortion parameter. We will exclude such a possibility for this study, assuming instead that the updated parameter values are a function of only the current parameter value and the persuasion actions.

Constraints (2e) – (2n) create the mappings, $T_{ijkk'}^+$ and $T_{ijkk'}^-$, from the new gain

and loss values, x_{ijk} , to their sorted, t_{ijk}^+ and t_{ijk}^- , values. Constraints (2e) and (2f) enforce a bijective mapping. Constraints (2g) and (2h) perform the mapping of outcomes for each decisionmaker, and Constraints (2m) and (2n) do likewise for probabilities. Constraints (2i) and (2j) ensure the values are in ascending order, and Constraints (2k) and (2l) enforce the positivity and negativity of gains and losses, respectively. The decision weights for gain and losses are calculated and enforced via Constraints (2o) – (2r) utilizing the Kahneman probability weighting functions specified in Equations (3a) and (3b). Constraint (2s) calculates the difference between prospect j and the preferred prospect q_i for each individual, utilizing the previously described decision weights and the Kahneman value functions listed in Equations (3c) and (3d). Constraints (2t) – (2v) ensure only one of the decision variables s_{ij}^{pos} or s_{ij}^{neg} assumes the entirety of this value, depending on the sign of the difference. Note that a value of zero for $z_{ij}^{q_i}$ is always feasible but, if properly constructed, the relation of the objective function to Constraint (2w) will induce a value of one whenever feasible. Constraint (2w) ensures an individual can only be labeled a supporter if he prefers prospect q_i to all of the other $n_i - 1$ prospects.

Persuasion actions are broadly defined in this formulation. They may be discrete or continuous, depending on the context. Moreover, the space A of feasible actions is assumed to be bounded. Likewise, the measure associated with the value of an outcome is not explicitly defined in this formulation. CPT usually assumes the measure is monetary in nature. Utilizing actuarial practices, monetary values can be applied to many non-monetary outcomes; however, these measures may still be ill-suited to describe some situations wherein persuasion is of interest. For example, in a presidential election, the monetary effects of a choice have value, but so do a wide array of other measures. Thus, it is probable that some situations may be better described by adapting a multi-attribute measure in the vein of Value-Focused Thinking (Keeney

and Keeney, 2009). Depending on the setting, these measures and their associated weights may be specific to individuals, making their estimation challenging. In such situations, the adaption of a discrete or continuous Likert Scale may be more effective. This course of action is promoted by Gass and Seiter (2015) who provided self-reporting scales (e.g., Likert Scales or Semantic Differential Scales) as an option to explicitly measure persuasion.

An interesting property of this formulation, in part driven by the restrictions on the parameter update functions f_{ijk} , g_{ijk} , and h_i^θ , is that there always exists a feasible solution. Whenever no attempt at persuasion is made, the parameter update functions are assumed to equal zero. Thus, the sorted outcomes t_{ijk}^+ and t_{ijk}^- are merely the $(\hat{y}_{ijk} - r_i)$ -values, listed in ascending order and separated by sign (i.e., positivity or negativity). Under such a null action by the persuader, the objective function value can be determined simply by calculating the prospect values for all individuals, counting those individuals who strictly prefer prospect q_i to all others, and substituting those values into the specific form of the objective function.

Finally, many situations may not allow for all prospect theory parameters, raw outcomes, and probabilities to be simultaneously affected. In such a scenario, the appropriate decision variables can be affixed as parameters and the presented formulation simplified. The malleability of outcomes such that rank-order is not preserved is perhaps the most complicating feature of this scenario. If rank-order can be assumed to be preserved, then the persuasion program can be greatly simplified.

Adaptations from Risk to Uncertainty.

To model an uncertain environment, we assume some probability-like measure is utilized to quantify uncertainty. Such an assumption accords with the belief-based model of decision under uncertainty by Fox and Tversky (1998) and is validated by the

work of Gonzalez and Wu (1999) and Kilka and Weber (2001). These authors quantified uncertainty using Support Theory developed by Tversky and Koehler (1994), wherein judged probabilities are measured by the support of a focal hypothesis compared to an alternative. That is,

$$P(A, B) = \frac{s(A)}{s(A) + s(B)},$$

wherein the function $s(\cdot)$ is characterized as the degree of support for a given hypothesis.

Under Support Theory, judged probabilities lose the property of extensionality. The probability of a disjunction of two independent and mutually exclusive events can no longer be assumed to be their sum because the implicit disjunction of a two events is assumed to be perceived as less likely than their explicit disjunctions. Namely, if hypothesis A is equal to the union of mutually exclusive hypotheses B and C , then A is an implicit disjunction, whereas $B \cup C$ is an explicit disjunction. In this setting, the authors propose $s(A) \leq s(B \cup C)$.

Therefore, to adapt our model to a belief-based approach using Support Theory, we introduce the sets C_{ij} and Ω_u .

$$C_{ij} : \text{Set of all } \Omega_u \text{ subsets of } K_{ij} \text{ such that } \{\Omega_u \in K_{ij} : 1 \leq |\Omega_u| \leq m_{ij}\}$$

The set C_{ij} is a collection of subsets of K_{ij} . Each of these subsets, Ω_u , represents a possible disjunction of the uncertain events in K_{ij} . CPT utilizes the assessed likelihood of these disjunctions to determine the decision weights. In the case of risk, these disjunctions are merely the sum of the component probabilities. However, as the assumption of extensionality fails to hold, the judged probability of each disjunction must be obtained. As such, the decision variables \hat{p}_{ijk} and p_{ijk} from the risk model

must be substituted with the following.

$\hat{p}_{ij\Omega_u}$: The judged probability that the disjunction of events in Ω_u occurs

$p_{ij\Omega_u}$: The judged probability that the disjunction of events in Ω_u occurs

after persuasion

To account for uncertainty in lieu of risk within the general model, one may respectively replace Constraints (2c) and (2m)–(2r) with Constraints (4a)– (4g) to leverage these Support Theory based, disjunction-oriented probabilities.

$$p_{ij\Omega_u} = \hat{p}_{ij\Omega_u} + f_{ij\Omega_u}(\mathbf{a}), \quad \forall i \in I, j \in J_i, \Omega_u \in C_{ij}, \quad (4a)$$

$$b_{i,j,k'}^+ = \sum_{\Omega_u: |\Omega_u|=m_{ij}-k'+1} p_{ij\Omega_u} \left[\prod_{l \in \Omega_u} \left(\sum_{l'=k'}^{m_{ij}} T_{ijll'}^+ \right) \right], \quad \forall i \in I, j \in J_i, k' \in K_{ij}, \quad (4b)$$

$$b_{i,j,k'}^- = \sum_{\Omega_u: |\Omega_u|=k'} p_{ij\Omega_u} \left[\prod_{l \in \Omega_u} \left(\sum_{l'=1}^{k'} T_{ijll'}^- \right) \right], \quad \forall i \in I, j \in J_i, k' \in K_{ij}, \quad (4c)$$

$$\pi_{ijm_{ij}}^+ = w^+(b_{ijm_{ij}}^+, \gamma_i), \quad \forall i \in I, j \in J_i, \quad (4d)$$

$$\pi_{ijk'}^+ = w^+(b_{ijk'}^+, \gamma_i) - w^+(b_{ij(k'+1)}^+, \gamma_i), \quad \forall i \in I, j \in J_i, k \neq m_{ij}, \quad (4e)$$

$$\pi_{ij,1}^- = w^-(b_{ij,1}^-, \delta_i), \quad \forall i \in I, j \in J_i, \quad (4f)$$

$$\pi_{ij,k'}^- = w^-(b_{ij,k'}^-, \delta_i) - w^-(b_{ij(k'-1)}^-, \delta_i), \quad \forall i \in I, j \in J_i, k' \neq 1, \quad (4g)$$

Constraint (4a) is a natural analog to Constraint (2c), except that a separate persuasion update function is required for each event disjunction. The same applies for Constraints (4d)–(4g). However, Constraints (4b) and (4c) are only required when a persuasion action can disrupt the rank-order of outcomes; they ensure that the proper disjunctions are assigned to the appropriate positions for calculation of the decision weights by utilizing the decision variables associated with the order mapping. Furthermore, if a measure of evidence strength can be ascertained for an uncertainty, Constraint (4a) can be written in terms of the persuasion action's effect on this metric.

Tversky and Koehler (1994) and Fox (1999) demonstrated that judged probabilities can be estimated by the direct assessment of strength of evidence. They show that

$$P(A, A^c) = \frac{s(A)}{s(A) + s(A^c)} = \frac{\hat{s}(A)^k}{\hat{s}(A)^k + \hat{s}(A^c)^k} = \frac{\hat{s}(A)^k}{\hat{s}(A)^k + w_{A^c} [\sum_{i=1}^n \hat{s}(B_i)^k]} \quad (5)$$

wherein $s(A)$ is the support of hypothesis A , $\hat{s}(A)$ is an evidence strength metric for A , events B_1 to B_n form the explicit disjunction comprising the implicit disjunction A^c , and the global weight w_{A^c} represents the ratio of support between these implicit and explicit disjunctions.

In such a situation, once values of k and w_{A^c} have been identified via regression analysis of data pertaining to observed behaviors and the effect of persuasion on $\hat{s}(A)$ ascertained, Constraint (4a) can be rewritten in a form similar to Equation (5). Such a functional relationship is highly beneficial, especially when the evidence strength metric is observable. In Section 2.4, we provide an example demonstrating this concept.

2.4 Example Applications

The estimation of all parameter update functions in a GPP formulation under either risk or uncertainty is feasible, but it is also expensive and dependent upon the specific persuasive action. Previous data collection efforts in human subject research are generally insufficient to simultaneously estimate all of these functional forms. The dynamic nature of risk attitude is alluded to in the literature, but few studies have translated these changes quantitatively into the complete CPT framework. In this section, we demonstrate the efficacy of our proposed model to specific contextual instances wherein a persuader can reasonably assume a subset of the decision variables

to be constant. By doing so, we demonstrate the flexibility of the GPP and provide motivation for estimating the dynamic effects of persuasion on CPT parameters via future human subject research.

The first application demonstrates the model under risk and considers groups of decisionmakers at an aggregate level. The persuader is able to alter the respective outcomes, probabilities, and reference points of two groups of decisionmakers. The probability weighting and utility functions are assumed to be constant. The remaining examples demonstrate the model's efficacy under uncertainty, with one modeling decisionmakers at an individual level and the other at the aggregate level. In the second of the three models, the persuader is able to affect judged probabilities for a set of individuals by strengthening or weakening the support for a hypothesis. Likewise, the uncertainty weighting function can be altered through either an emotional appeal or by affecting source preference. The third model assumes only the prospects faced by demographic groups can be affected and the remaining elements are constant. All instances are solved utilizing the global solver BARON on an HP ZBook equipped with a 2.70 GHz Intel i7-4800MQ processor and 32GB of RAM. Our selection of a commercial solver is informed by the results of Caballero et al. (2018) on a similar mathematical program having sorting constraints coupled with functions that use decision variables in both the base and exponent of an expression. For further information regarding BARON, we refer the interested reader to selected research conducted to develop and improve the solver's algorithms, as conducted by Ryoo and Sahinidis (1995, 1996); Tawarmalani and Sahinidis (2004, 2005); Khajavirad and Sahinidis (2013); Sahinidis (2018); and The Optimization Firm (2019).

Influence under Risk: Insurance Policies.

Insurance is a popular application area for Cumulative Prospect Theory. In an extensive review of its descriptive capability, Camerer (2000) discussed how prospect theory explains phenomena with respect to automobile, health, and telephone wire insurance. Given this interest in the literature, our first application of the GPP relates to a home insurance company designing policies to maximize their expected profit.

For this use case, we make the following assumptions. The insurance provider is developing a business model with “no hassle” policies. All homes of a given type are provided a ready-made plan with set premiums and deductibles. Two aggregate groups of homeowners of home types $i \in I = \{1, 2\}$ each have two lotteries $j \in J = \{Buy Policy, Opt Out\} = \{1, 2\}$ from which to chose. It is assumed that any one of four events can occur to a home, where $K_{ij} = \{Fire, Hail, Wind, Nothing\} = \{1, 2, 3, 4\}$ for every homeowner group and prospect. Only one event is able to occur in a given payment period. A decision tree for homeowner group i based on the insurer’s decision variables can be seen in Figure 3. Of note, we assume $\hat{y}_{i1k} = 0$ and $\hat{p}_{i1k} = \hat{p}_{i2k} = \hat{p}_{ik}, \forall i \in I, k \in K_{1i}$, where \hat{p}_{ik} is defined for notational convenience for this use case only.

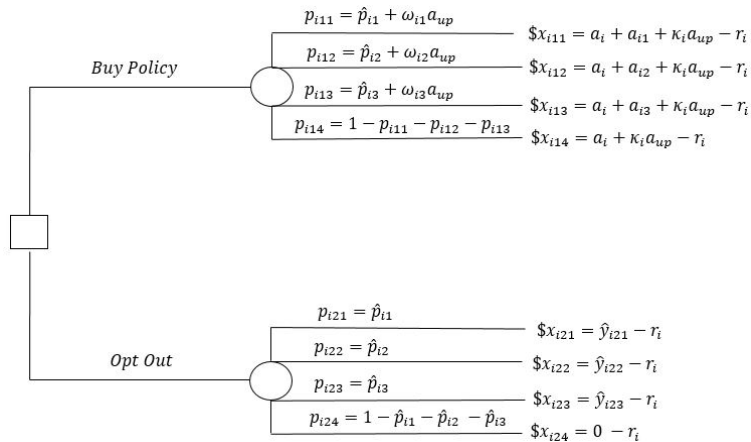


Figure 3. Customer Decision Tree for Purchasing Insurance Policy

This example is designed to illustrate our model's applicability towards risk. Therefore, in the vein of an actuarial example, we make the assumption that frequentest probabilities and associated monetary damages of an insurable event are known (or can be reasonably estimated) and are utilized to inform both customer and insurer decisions. The probability of any given event is assumed to be disparate for each home type. The insurance provider is assumed to use aggregate data in setting prices. As such, the original probability weighting function proposed by Tversky and Kahneman (1992) is adopted, in addition to their proposed (power) utility function and the parameters associated with these functions (e.g., loss aversion parameter, utility curvature, etc.).

Coinciding with the emergence of separate wind/hail deductibles in the U.S. market, the insurer's policies incorporate both fixed rate premiums (a_i) and event-based deductibles (a_{ik}). In addition to setting the premium and deductible amount for each policy, the insurance provider faces two additional decisions: (1) whether to require all policy holders to upgrade their roof to a cutting-edge material costing κ_i that affects the hazard probability by ω_{ik} ; and (2) whether to implement a controversial advertising campaign of cost $\sigma = \$500$ to emphasize household hazards that could either impel insurance purchase or harden resolve against it, depending on the decisionmaker. These binary decisions are respectively represented as a_{up} and a_{ad} . Additionally, a binary variable ν_{ik} is introduced to determine whether the damage incurred exceeds the deductible.

The formulation of this instance utilizes the majority of the constraints in the proposed GPP under risk. However, there do exist a few modifications, as the parameters within both the probability weighting function and utility functions are assumed to be constant. The insurance provider operates via Expected Utility as a risk neutral decisionmaker and optimizes its profit using the following mathematical

program, assuming customers purchase insurance in accordance to CPT. We note that the ‘‘Opt Out’’ prospect preserves rank order, and the base probabilities are not affected. Therefore, we need only consider the case where $j = 1$ for many constraints.

$$\max -\sigma a_{ad} + \sum_{i \in I} \Phi_i \left(\sum_{k \in K_{i1}} \nu_{ik} p_{i1k} (-\hat{y}_{i2k} + a_{ik} + a_i) \right) \quad (6a)$$

subject to Constraints (2e)–(2v), (2y)–(2ab),

$$\nu_{ik} a_{ik} \leq \hat{y}_{i2k}, \quad \forall i \in I, k \in K_{i1}, \quad (6b)$$

$$a_{ik} \geq (1 - \nu_{ik}) \hat{y}_{i2k}, \quad \forall i \in I, k \in K_{i1}, \quad (6c)$$

$$x_{i1k} = a_i + a_{ik} + \kappa_i a_{up} - r_i, \quad \forall i \in I, k \in K_{i1}, \quad (6d)$$

$$x_{i2k} = \hat{y}_{i2k} - r_i, \quad \forall i \in I, k \in K_{i2}, \quad (6e)$$

$$p_{i1k} = a_i + \omega_{ik} a_{up}, \quad \forall i \in I, k \in K_{i1}, \quad (6f)$$

$$r_i = \hat{r}_i + \rho_i a_{ad}, \quad \forall i \in I, \quad (6g)$$

$$\Phi_i = z_{i2}^1, \quad \forall i \in I, \quad (6h)$$

$$\nu_{ik} \in \{0, 1\}, \quad \forall i \in I, k \in K_{i1}, \quad (6i)$$

$$0 \leq a_{1k} \leq 5,000, \quad \forall k \in K_{11}, \quad (6j)$$

$$0 \leq a_{2k} \leq 10,000, \quad \forall k \in K_{21}, \quad (6k)$$

$$0 \leq a_1 \leq 2,000, \quad (6l)$$

$$0 \leq a_2 \leq 4,000. \quad (6m)$$

Table 1 presents the correspondence of constraints between this particular formulation and the general GPP formulation set forth in Section 2.3. The objective function (6a) calculates the total expected profit obtained by the insurer, deducting expected expenditures from expected revenue. Constraints (6b) and (6c) determine whether damage exceeds the set event-based deductible, whereas Constraints (6f)–(6g) update the outcome values, probabilities, and reference points specific to this context. Likewise, in this setting there is only one alternative prospect, implying Φ_i can be determined directly via Constraint (6h). Although the formulation can be fur-

ther simplified for this instance by replacing Φ_i with z_{i2}^1 throughout the formulation and eliminating Constraint (4i), we refrain from doing so to maintain consistency of terminology.

Table 1. Correspondence of Instance-specific and GPP constraints

Instance	GPP	Remarks
(6a)	(2a)	
(2e)-(2v), (2y)-(2ab)	As presented	
(6b), (6c), (6i)	N/A	These constraints are used to compute the instance-specific objective function
(6d)-(6g)	(2b)-(2d)	
(6h)	(2w)	The GPP constraint (2w) can be simplified because there are only two prospects
(6j) - (6m)	(2x)	

The values for all parameters for each homeowner group are presented in Table 2. Each homeowner may have disparate reference points pertaining to the amount they expect to pay for home repairs. Thus, we see homeowners of Type 1 have a reference point of $-\$1,000$, whereas homeowners of Type 2 have a reference point of $\$0$. The cost of upgrading a roof (κ_i) must be born by the policy holder and will reduce the hazard of some insurable events but may increase it for others. Specifically, we assume the upgraded roof is more resilient to wind and hail but more susceptible to fire damage, as it is made of a more flammable material. As such, we provide the corresponding p_{ik} and ω_{ik} -values in Table 2. Likewise, the advertising campaign is anticipated to affect the reference point for customers, as represented by the values of ρ_i . Ceilings of $\$2,000$ and $\$5,000$ are respectively imposed on the premium and deductible amounts for Home Type 1, and ceilings of $\$4,000$ and $\$10,000$ are imposed for these values on Home Type 2. These limits are applied to ensure the plans remain economically feasible for the socio-economic condition of the region and bound both a_i and a_{ik} from above. It is possible for such constraints to make it unprofitable to offer insurance to a specified home type. In this situation, additional modifications would be required to the GPP. However, this example has been constructed such that

the insurer can profit by offering insurance within the socio-economic limits. Finally, a discriminable factor of $\Delta = 5$ is adopted for ascertaining whether a homeowner prefers to buy the policy.

Table 2. Parameter Values Associated with Insurance Persuasion Program

Parameter	Home Type 1	Home Type 2
Value	\$100,000	\$200,000
Reference Point (\hat{r}_i)	-\$1,000	\$0
Fire Probability (p_{i1})	0.005	0.001
Hail Probability (p_{i2})	0.02	0.03
Wind Probability (p_{i3})	0.01	0.01
Fire Damage (\hat{y}_{i21})	-\$70,000	-\$160,000
Hail Damage (\hat{y}_{i22})	-\$20,000	-\$30,000
Wind Damage (\hat{y}_{i23})	-\$11,000	-\$18,000
Upgrade Fire Probability Effect (ω_{i1})	0.001	0.001
Upgrade Hail Probability Effect (ω_{i2})	-0.015	-0.01
Upgrade Wind Probability Effect (ω_{i3})	-0.0075	-0.005
Advertising Reference Point Effect (ρ_i)	\$500	-\$2,000
Upgrade Homeowner Cost (κ_i)	-\$2,000	-\$4,000

The commercial solver BARON was allotted five hours of computational effort to identify a global optimal solution. The time limit was reached prior to convergence, but the reported maximum amount the insurer can earn per period for a set of these two home types is \$3,095.26. Therefore, optimality of this solution is not guaranteed. However, it does provide the insurer with a lower bound on the expected profit that can be attained. This level of profit is obtained by foregoing the advertisement but instituting the mandatory roof upgrade policy. The deductibles for the \$100,000 home are listed as $\{\$4,915.52; \$4,999.50; \$4,919.150\}$ and for the \$200,000 home as $\{\$0.67; \$0.70; \$0.53\}$, and their respective premiums are \$1,999.93 and \$2,586.00.

In this solution, owners of both home types prefer to buy their offered policy. However, the policies exhibit different behavior with regard to their deductibles; Home Type 1 is required to pay substantial deductibles, whereas Home Type 2 has virtually no deductible for any event. This behavior is due to the socio-economic constraints placed on these values, in addition to the relation of the premium and the deductible to the profit function. Positive profit is only obtained by increasing the premium,

whereas an increase in the deductible merely curtails an expenditure. As such, if able to do so, the insurer prefers all payment in the form of the premium and only resorts to increasing the deductible values when the premium is near the homeowner's socio-economic limit. This behavior is appropriate for our simplified model because we ignore moral hazards and extraneous claims, two commonly cited reasons for implementing a deductible. Were these considerations desired to be modeled, an alternative lower bound on the deductibles could be instituted, but doing so would likely lessen the expected profit through an associated premium reduction.

Influence under Uncertainty with Evidence Strength Metric: Defending a Client.

A defense attorney's goal is to convince a jury of their client's innocence. The jurors, for their part, are asked to listen to both the prosecution's and the defense's arguments to determine whether the accused is guilty beyond a reasonable doubt. Diamond (2003) defended the jury as a decisionmaker in response to criticism of the U.S. Constitution's Seventh Amendment and concluded that juries are competent, albeit imperfect decisionmakers due to their human nature. In this example, we show how an attorney can use the tools of decision under uncertainty to leverage the human nature of the jury to their benefit.

In this context, each juror's decision can be viewed as a decision under uncertainty, as depicted in Figure 4. Jurors are not provided with the probabilities of guilt but must infer them through the arguments presented in court. A verdict of guilty is preferred if a juror assesses the probability of guilt to exceed a threshold of reasonable doubt. This threshold of reasonable doubt is determined by the utilities associated with either a correct or an incorrect verdict.

Consider a defense attorney in a criminal trial trying to achieve an acquittal for

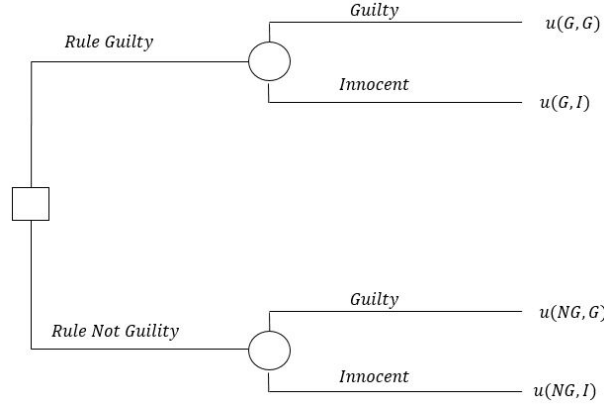


Figure 4. Juror Decision Tree: Two Available Verdicts

his client (i.e., not a hung jury or a mistrial). Two of the twelve jurors are perceived through *voir dire* to be of such strong personality, it is determined that, if they can be convinced of acquittal, the rest of the jury will follow. We assume the prosecution has just rested their case, and the defense attorney is deciding how to respond. Utilizing the 0–100 suspicion scale from Tversky and Koehler (1994), the defense attorney estimates the prosecution was very effective and that each of the two jurors of interest views his client as a 95 (i.e., most likely guilty). However, through cross examination, the defense attorney was able to cast suspicion on three other individuals such that their respective values on the suspicion scale are 25, 20, and 15. Even though the three other individuals are not on trial, if a juror has high suspicion values for the group compared to the person on trial, then reasonable doubt will exist. We define event A as the defendant having committed the crime (i.e., defendant is guilty) and events B , C , and D as one of the three individuals referenced by the defense having committed the crime (i.e., the defendant is not guilty).

In this setting, it can be observed that $I = \{Juror\ 1, Juror\ 2\} = \{1, 2\}$, $J_i = \{Rule\ Guilty, Rule\ Not\ Guilty\} = \{1, 2\} \forall i$, and $K_{ij} = \{Guilty, Innocent\} = \{1, 2\} \forall i \in I, j \in J$. Should we assume the example estimates from Fox (1999) of $k = 2.2$ and $w_{A^c} = 0.65$ to be applicable and static, the judged probability that juror i assigns

to the defendant of being guilty (i.e., $\hat{p}_{i11} = P(A, A^c) = \hat{p}_{i21}$) can be calculated via Equation (7). We further note that the support assigned to the defendant not committing the crime is assumed to equal the weighted support that one of the other individuals did. However, it is assumed the jurors do not unpack A^c on their own (i.e., with additional hypotheses regarding who, specifically, might be guilty, if not A); their task is to consider the collective set of alternatives, $A^c = B \cup C \cup D$. Due to binary complementarity, it can also be observed that both \hat{p}_{i12} and \hat{p}_{i22} are equal to $P_i(A^c, A) = 1 - \hat{p}_{i11} = 1 - \hat{p}_{i21}$.

$$\hat{p}_{ij1} = \frac{\hat{s}_i(A)^k}{\hat{s}_i(A)^k + w_{A^c}[\hat{s}_i(B)^k + \hat{s}_i(C)^k + \hat{s}_i(D)^k]} \quad (7)$$

A common reference point of zero is assumed for all jurors along with a Δ -parameter of zero. Jurors are further assumed to have Linear-in-Log-Odds probability weighting functions with parameter values corresponding to experiments conducted by Kilka and Weber (2001), which are identical for gains and losses such that, for each juror i , we have $\hat{\delta}_i = \{1.096, 0.953\}$ and $\hat{\gamma} = \{0.489, 0.415\}$. The general form of the Linear-in-Log-Odds function is stated in Equation (8).

$$w^+(p) = w^-(p) = \frac{\delta^+ p^{\gamma^+}}{\delta^+ p^{\gamma^+} + (1-p)^{\gamma^+}}. \quad (8)$$

We assume these estimates correspond with familiar and unfamiliar sources of uncertainty from the Kilka and Weber (2001) experiments. As such, we assume Juror 1 is familiar with the underlying uncertainty whereas Juror 2 is not. The reference point adjusted utility for a given outcome is assumed to be fixed such that a correct verdict yields $u(G, G) = 1 = u(NG, I)$, an incorrect guilty verdict yields $u(G, I) = -5$, and an incorrect not guilty verdict yields $u(NG, G) = -4$. Therefore, the persuasion

action preserves rank order of outcomes.

The prosecutor has seven themes to utilize in crafting a defense argument, and any combination of these themes may be employed. A binary decision variable a_l is associated with each of the seven themes, each of which has the potential to increase or decrease the suspicion for a given individual, with the specific effect dependent upon the juror. For instance, one potential theme may be to insinuate that a corrupt government official is the real culprit. From jury selection, the defense attorney may know that one juror has experienced mistreatment from authority figures and may identify personally with such an argument. However, if the second juror is a lifetime civil servant, such a theme may not resonate with them. Support for some hypothesis $E \in \{A, B, C, D\}$ after persuasion is denoted as $\bar{s}_i(E)$.

We also assume the themes utilized can affect the probability weighting function of either of the two jurors. This effect may result from a theme's emotional content or the jurors' respective familiarity with the source of uncertainty. For example, if the underlying suspicion of the defendant is rooted in ethical finance practices, a selected theme may include educational aspects relating to the nuances of this profession. Should one juror be unfamiliar with these nuances, this theme may help reduce the adverse source preference effect. For this illustrative example, the quantitative impact of each theme against each hypothesis and juror's probability weighting function is shown in Table 3, wherein τ_{iEl} represents the effect of Theme l on the support for hypothesis E for juror i .

To model these effects, new binary decision variables a_l must be introduced, which equal one if Theme l is utilized and zero otherwise. The GPP under uncertainty can be updated as follows. The objective function (9a) assumes the defense attorney seeks to select a combination of themes to maximize the number of jurors favoring a *Not Guilty* verdict prior to deliberation. The support update function for each

Table 3. Parameter Update Coefficients Associated with Defense Themes

Effect	Theme(<i>l</i>)						
	1	2	3	4	5	6	7
γ_1	0	0	-0.088	0	-0.015	0	-0.05
δ_1	0	0	-0.1	0	-0.025	0	-0.05
τ_{1A}	-10	-5	5	-5	-5	-15	0
τ_{1B}	10	-5	10	-5	5	-10	5
τ_{1C}	0	10	-10	20	-10	30	0
τ_{1D}	0	10	-10	20	-10	30	0
γ_2	0.08	0.04	-0.15	0	0	-0.01	0
δ_2	0.1	0.05	-0.002	0	0	-0.01	0
τ_{2A}	-5	10	-5	-5	-5	-15	0
τ_{2B}	10	5	-10	5	5	-10	5
τ_{2C}	0	5	-5	10	-15	50	0
τ_{2D}	10	0	0	0	-5	-5	10

hypothesis can be seen in Equation (9b), and the probability update functions are of the form shown in Equations (9c) and (9d). Based on the context of our problem, there are only two outcomes for each prospect: one outcome represents a gain and the other outcome represents a loss. As rank order is preserved, the Prospect Order Constraint Set is not required and, when coding the instance, we can take the b_{ijk}^{\pm} - and t_{ijk}^{\pm} -variables as their appropriate p_{ijk} - and x_{ijk} -counterparts (e.g., $p_{111} = b_{112}^+$ and $p_{121} = b_{121}^-$). Binary complementarity also ensures the sum of judged probabilities equals one, in accordance with Constraint (9d). For reference, Table 4 presents the correspondence of constraints between the particular and general formulations for this instance of the GPP.

$$\max \sum_{i \in I} \Phi_i \tag{9a}$$

subject to Constraints (4d)–(4g), (2s)–(2aa),

$$\bar{s}_i(E) = \hat{s}_i(E) + \sum_{l=1}^7 \tau_{iEl} a_l, \quad \forall i \in I, E \in \{A, B, C, D\}, \tag{9b}$$

$$p_{ij1} = \frac{\bar{s}_i(A)^k}{\bar{s}_i(A)^k + w_{Ac} [\bar{s}_i(B)^k + \bar{s}_i(C)^k + \bar{s}_i(D)^k]}, \quad \forall i \in I, j \in J_i, \tag{9c}$$

$$p_{ij2} = 1 - p_{ij1}, \quad \forall i \in I, j \in J_i. \tag{9d}$$

Solving the adapted persuasion program in BARON, an optimal solution is attained in 0.08 seconds of computational effort. The defense attorney is able to persuade both jurors of his client’s innocence by utilizing the Themes 1, 2, 4, 5, and 6. The respective difference in prospect values is 1.37 and 0.026 in favor of acquittal for Jurors 1 and 2.

Table 4. Correspondence of Instance-specific and GPP constraints

Instance	GPP	Remarks
(9a)	(2a)	
(4d)-(4g), (2s)-(2aa)	As presented	
(9b)	N/A	Instance-specific constraint used to relate the evidence strength metric to the judged probability value
(9c)-(9d)	(4a)	

However, we note that there are alternative optimal solutions in this setting. For instance, if the previously reported optimal is disallowed via a new constraint, BARON reports the combination of Themes 4, 6, and 7 as an optimal solution having the same objective function value. Using this combination of themes, the respective difference between acquittal and conviction is 0.187 and 0.29 for Jurors 1 and 2. To distinguish between alternative optimal solutions, additional manpower, budgetary, or strategic considerations can be included via a subset of new constraints. For example, the number of themes to be utilized may be limited by the number of paralegals available for research. In turn, the number of paralegals may be affected by the law firm’s budget and total caseload. The inclusion of this dynamic in the constraints enables the attorney to select an optimal defense strategy while considering concomitant factors.

Influence under Uncertainty without Evidence Strength Metric: Pension Enrollment.

In our final example, we consider a public agency reconstructing its pension program for government employees (e.g., see Herald, 2016). The public agency has been operating a legacy, pension-only system that is no longer affordable. The legacy pension system guarantees an annual pension of 50% of salary at retirement (μ_{old}) to employees who stayed with the company for 25 years. Income during employment with the company is based on a time-of-employment stair schedule and is readily predictable. Alternatively, the new mixed pension system reduces the annuity percentage to 40% of the salary at retirement (μ_{new}), but it includes an additional retirement benefit based on employee contributions to a stock market fund over their duration of employment. We assume that, if the stock market has increased in a given year, an additional τ is provided to the retiree whereas, if the market has declined, a smaller value of ρ is provided. All new hires will be switched to the new, mixed-pension system, but the agency wishes to reduce future payroll costs by incentivizing current employees to switch to the new retirement plan. Assuming a choice is made contingent upon expected annual retirement income, the decision tree faced by these employees can be seen in Figure 5. All outcome components of the decision tree are fixed except for the continuous variable a . For this scenario, the variable a represents the amount (in dollars) of an additional incentive called continuation pay which will be provided to any employee for switching to the new retirement system. However, this continuation pay can only be retained if the member chooses to remain employed with the agency until retirement.

To avoid complications with the time value of money, we assume all values have been adjusted to current year dollars. The agency is considering four demographics it would like to impel to switch into the new pension system. Each demographic has

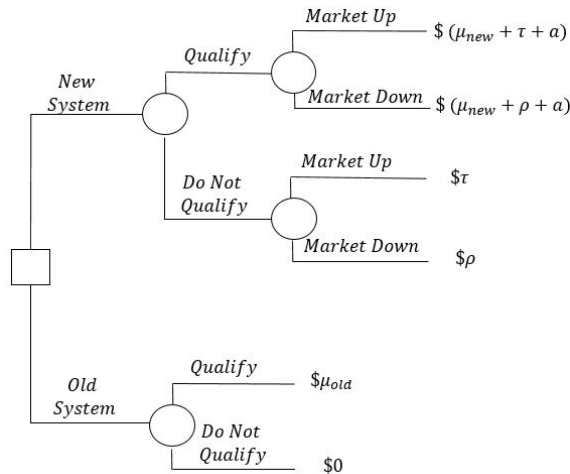


Figure 5. Pension Enrollment Decision Tree

varying values of μ_{old} , μ_{new} , τ , and ρ parameters based on their projected income at retirement, as reported in Table 5. Each demographic is assumed to utilize the Cumulative Prospect Theory functions and parameters from Tversky and Kahneman (1992), with a reference point of zero. The parameter values for these CPT functions are not malleable.

The public agency is only able to affect the outcomes depending on the value of a . As such, the rank order of outcomes in this setting is not permutable. We recognize that the assumption of a reference point equal to zero may be subject to alternative views; one could argue that an individual's reference is their expected retirement income. However, we justify herein the reference point of zero for both tractability and demonstration purposes.

A compressed version of the decision tree for use in our persuasion program is illustrated in Figure 6. As rank order is preserved, we assume the public agency has polled its employees over the appropriate event disjunctions utilized in the probability weighting functions. In this way, the agency directly assesses the judged probability of its constituents for each event and utilizes them in the appropriate position to calculate the decision weights. These values are also listed in Table 5. It is assumed

Table 5. Parameter Update Coefficients Associated with Pension Persuasion

Parameters	Demographic 1	Demographic 2	Demographic 3	Demographic 4
μ_{old}	\$20,000	\$30,000	\$40,000	\$50,000
μ_{new}	\$16,000	\$24,000	\$32,000	\$40,000
τ	\$3,000	\$4,500	\$6,000	\$7,500
ρ	\$400	\$600	\$800	\$1,000
\hat{p}_{i11}	0.35	0.1	0.3	0.4
\hat{p}_{i12}	0.2	0.2	0.25	0.55
\hat{p}_{i13}	0.3	0.3	0.3	0.05
\hat{p}_{i14}	0.25	0.5	0.25	0.075
$\hat{p}_{i1(1,2)}$	0.5	0.2	0.45	0.85
$\hat{p}_{i1(1,2,3)}$	0.8	0.6	0.75	0.95
\hat{p}_{i21}	0.6	0.3	0.5	0.9
\hat{p}_{i22}	0.4	0.7	0.5	0.1

that the event space associated with the uncertainty of each prospect cannot be further subdivided. As such $\hat{p}_{i1(1,2,3,4)} = \hat{p}_{i2(1,2)} = 1, \forall i \in I$, per Support Theory axioms.

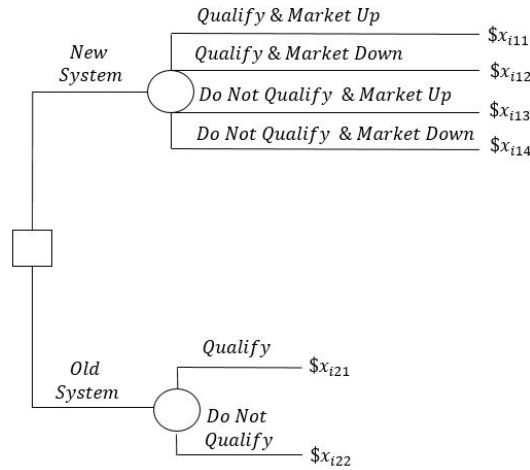


Figure 6. Compressed Pension Enrollment Decision Tree

The following adaption of the GPP under uncertainty allows the agency to determine the minimum continuation pay it must provide to convince all demographics to switch to the new retirement system. Table 6 presents the correspondence of constraints between the particular and general formulations for this instance. As with the example presented in Section 2.4, note that the assumptions in this instance allow for the exclusion of selected GPP constraints. The outcome update functions determining x_{ijk} -values correspond to the events depicted in Figure 5. It is assumed

that, for each dollar a increases, each probability associated with the new retirement system is positively incremented by, e.g., a value of 2.5×10^{-6} . Each \hat{y}_{ijk} equals the fixed elements of the associated outcome value (e.g., $\hat{y}_{i11} = \mu_{new} + \tau$).

$$\min a \tag{10a}$$

subject to Constraints (4d) – (4g) , (2s)–(2aa),

$$x_{i1k} = \hat{y}_{ijk} + a, \quad \forall i \in I, k \in \{1, 2\}, \tag{10b}$$

$$x_{i1k} = \hat{y}_{ijk}, \quad \forall i \in I, k \in \{3, 4\}, \tag{10c}$$

$$x_{i2k} = \hat{y}_{ijk}, \quad \forall i \in I, k \in K_{i2}, \tag{10d}$$

$$p_{i1\Omega_u} = \hat{p}_{i1\Omega_u} + a * (2.5 \times 10^{-6}), \quad \forall i \in I, \Omega_u \in C_{i1}, \tag{10e}$$

$$p_{i2\Omega_u} = \hat{p}_{i2\Omega_u}, \quad \forall i \in I, \Omega_u \in C_{i2}, \tag{10f}$$

$$\sum_{i \in I} \Phi_i = 4. \tag{10g}$$

Utilizing this information and a value of $\Delta = 0$, the resulting persuasion program is solved optimally with 0.69 seconds of computational effort by the commercial solver BARON. A minimum continuation pay of \$5,950 is required to impel the targeted demographics to switch pension programs. Each demographic favors switching to the new pension system by a margin of \$525.53, \$515.83, \$615.80, \$0.20, respectively, in CPT-calculated value.

Table 6. Correspondence of Instance-specific and GPP constraints

Instance	GPP	Remarks
(10a)	(2a)	
(4d)-(4g), (2s)-(2aa)	As presented	
(10b)-(10f)	(2b)-(2c)	
(10g)	N/A	Instance-specific constraint ensuring all decisionmakers are persuaded

Since Demographic 4 finds the two deals very similar, the agency may need to consider resolving with a higher Δ -value if it wants to ensure all demographics switch to the new system. Assuming the agency has determined the scaled continuation

pay cannot exceed \$7,000 without reversing the cost savings of the mixed pension system, the program is resolved with $\Delta = 200$. The optimal solution indicates all demographics will enroll in the mixed pension system with a minimum continuation pay of \$6,924 with margins similar to those aforementioned. This solution provides more assurance on employee behavior, but may reduce agency cost savings. Thus, in making its final decision on the continuation pay value, agency leadership must balance this risk versus reward dynamic, potentially by utilizing traditional Decision Analysis techniques.

2.5 Discussion and Conclusions

The concept of influencing an individual to adopt a preferred course of action is an ancient concept. The Ancient Greeks even personified persuasion in the goddess Peitho (Theoi Project, 2017). It is also nearly ubiquitous. Commercial marketing campaigns are persuasion operations wherein the desired prospect is a customer buying a given product. Political campaigns are persuasion operations seeking to garner votes for a particular candidate. Financial regulatory actions are persuasion operations engaging corporations and individuals, both of which are susceptible to emotional stimuli (Fairchild, 2014), with the objective of ensuring market confidence, financial stability, and consumer protection (Financial Services Authority, 2018). Additional application areas include management, deterrence, lobbying legislative bodies, counter-terrorism, social marketing, and many other forms of interpersonal communication.

A new class of decision problems has been introduced to model this human behavior. We formulated a model rooted in fundamental behavioral economic concepts such that a persuader can select an optimal strategy to meet a desired goal, subject to selected restrictions. By doing so, we fill a gap in the Behavioral Operations Re-

search literature by expanding the stream of research Franco and Hämäläinen (2016) described as “[modeling] human behavior in complex settings” from a descriptive to a prescriptive setting, and thereby strengthen the research stream’s connection to the core of Operations Research study. Likewise, a similar observation can be made when considering our work from the broader perspective of behavioral science; that is, it leverages descriptive theories of choice to inform an agent’s optimal interaction with a population. The models introduced in this research are unique in their ability to address a variety of forms of persuasion under conditions of either risk or uncertainty. It compliments psychological persuasion models (e.g., the Elaboration Likelihood Method) that focus on how a persuasive message is processed by examining how a *persuasion campaign* can be designed to achieve desired effects.

However, our models are limited by the lack of knowledge surrounding the functional mappings of persuasion actions to their effect on Cumulative Prospect Theory functional parameters and reference points. Future empirical studies in Behavioral Operations Research and other, related disciplines are ultimately required to enable the full potential of these models. Of particular interest are psychophysiological studies (e.g., Leppänen et al., 2018) that capture the effect of emotion. Given the potential effects of culture, background, age, and other factors to affect emotional responses to persuasive stimuli, such experiments may requisite a significant number of control variables and prove to be challenging-but-fruitful research endeavors. In the absence of such data to inform regression of the parameter update functions, the models set forth herein provide a framework to adopt for influence campaigns but would necessarily utilize less quantitative estimation methods. Alternatively, practical implementation under current conditions could be facilitated by other robust decisionmaking methodologies (e.g., Lempert et al., 2006).

Nonetheless, some of the first advances have already been shared in the litera-

ture to develop these functional mappings with a view towards broader application. Booij et al. (2010) provided a description of the risk attitudes of a variety of demographics in the Netherlands by characterizing the expected value of their Cumulative Prospect Theory parameters. The respective characterization of external action on the probability weighting function parameters by Campos-Vazquez and Cuijly (2014) and Schulreich et al. (2014) provided the first empirical estimates of a probability weighting update function. Future studies combining these methodologies can be utilized to generate the data required to fit the parameter update functions related to a specific persuasion action.

Other compelling areas for future research relate to the study of heuristic methods applied to the GPP, and to the automated discovery of parametric input values. As illustrated in Section 2.4, instances of our model can be constructed such that they are difficult to solve quickly to optimality, even with a global solver. It is a worthy research endeavor to explore under what conditions this behavior emerges and to develop methods for its mitigation. The use of such methods or, alternatively, the heuristic use of a global solver via the imposition of a time limit, are both promising strategies to approach large persuasion problems in practice. Moreover, in a manner analogous to that discussed by White et al. (2016), it may be possible to leverage big data and/or sentiment analysis to describe changes in judged probabilities. Social media provides a platform to either solicit or infer an individual's evaluation of strength for a given hypothesis and, if defined properly, their assessment of probability. If such research is successful, the persuasion programs modeled herein have the potential to shift the paradigm of strategic decisions regarding influence.

III. Informing National Security Policy by Modeling Adversarial Inducement and its Governance

Abstract

The distinction between peace and conflict in contemporary international relations is no longer well-defined. Leveraging modern technology, hostile action below the threshold of war has become increasingly effective. The objective of such aggression is often the influence of opinions, emotions, and, ultimately, the decisions of a nation's citizenry. This work presents two new game theoretic frameworks, denoted as *prospect games* and *regulated prospect games*, to inform defensive policy against these threats. These frameworks respectively model (a) the interactions of competing entities influencing a populace and (b) the preemptive actions of a regulating agent to alter such a framework. Prospect games and regulated prospect games are designed to be adaptable, depending on the assumed nature of persuaders' interactions and their rationality. The contributions herein are a modeling framework for competitive influence operations under a common set of assumptions, model variants that respectively correspond to scenario-specific modifications of selected assumptions, the illustration of practical solution methods for the suite of models, and a demonstration on a representative scenario with the ultimate goal of providing a quantifiable, tractable, and rigorous framework upon which national policies defending against competitive influence can be identified.

3.1 Introduction

The role of a nation's citizenry in models of interstate conflict is often captured via the concept of *audience costs*. Introduced by Fearon (1994), audience costs characterize the punishment levied by the population of a state upon its leader for backing

down in a crisis escalation situation. This concept has been extended and tested by numerous authors and is well-studied in the literature (e.g., Weeks, 2008; Moon and Souva, 2016; Chiozza, 2017; Tomz et al., 2018). However, this research thread tends to focus on the military instrument of national power and implies a binary relationship between war and peace.

Conversely, the United States Joint Chiefs of Staff (U.S. Joint Chiefs of Staff, 2018) underscore that multiple instruments of national power (e.g., diplomatic, informational, military, economic) can be coordinated and applied to achieve strategic objectives. Chinese and Russian strategists place similar emphasis on coordinating instruments of national power in their respective doctrines of Unrestricted Warfare and New Generation Warfare.

Moreover, “[conventional] Western concepts of war are incompatible and fundamentally misaligned with the realities of conflict in the twenty-first century” (Stowell, 2018). Empowered by modern technology, emerging national defense strategies, often described by the term hybrid warfare, frequently utilize tactics stopping short of conventional war. The result is what NATO Secretary General Jens Stoltenberg classified as a “new and more demanding security environment where also there is a more blurred line between peace and war” (Woody, 2018b). This perspective is further elucidated in the 2018 U.S. National Defense Strategy (U.S. Department of Defense, 2018) wherein interstate strategic competition of a multipolar, multi-domain, and ambiguous nature is described as the primary concern to U.S. national security.

In such strategic competition, a nation’s citizenry may be explicitly targeted, potentially by multiple instruments of power, often with the objective of influencing their support for (or against) some prospect. This type of targeting can be witnessed in recent Russian influence efforts in Ukraine (Stowell, 2018), the United States (DHS and FBI, 2016), and the Balkans (Krastev, 2019). Moreover, as witnessed by the

actions of China and North Korea in the 2016 American presidential election, multiple adversaries may compete simultaneously over the public's support (Silverstein, 2017; Dupuy, 2018).

As such, in 21st Century interstate competition, the role of a nation's citizenry transcends the traditional concept of audience costs. In fact, as the decision-making of a nation's citizenry may be the object of competition, it is necessary to model their framework for choices with high granularity.

From an offensive influence perspective, the persuasion programs set forth by Caballero et al. (2018) can be leveraged to generate optimal strategies in a leader-followers game. Within a persuasion program, an external entity known as the *persuader* acts to impel a certain decision (or certain respective decisions, in the case of persuading a group of decisionmakers) by affecting (1) the underlying risk or uncertainty, (2) subjective beliefs regarding payoffs, or (3) cognitive evaluation of the set of prospects. Conceptually, their framework can be viewed as an adversarial Decision Analysis problem wherein the persuader can affect the structure of an individual's decision tree (e.g., see Figure 7), and the Cumulative Prospect Theory (CPT) parameters utilized in its evaluation (Tversky and Kahneman, 1992). In this manner, the persuader and the decisionmakers adopt the roles of the leader and the followers, respectively, in a bilevel programming context. However, the assumption of boundedly rational followers via the use of CPT allows for the problem's reformulation to a single-level form.

The research herein extends their methodology to a defensive setting wherein a third party, henceforth referred to as a *regulator*, is responsible for ensuring the actions of multiple persuaders fall within certain confines of acceptable behavior. We do so by introducing *prospect games* and *regulated prospect games*. A prospect game can be viewed as a persuasion program having multiple competing persuaders such that each




Figure 7. Example Decision Tree

has a finite number of available actions, whereas a regulated prospect game considers how a regulator would optimally interact with the persuaders in this setting. That is, a regulated prospect game is a prospect game acted upon first by a rational external entity (i.e., one who abides by Expected Utility axioms).

To inform the regulator’s decision on whether to intervene, the prospect game should first be solved to identify outcomes in the absence of regulator action. Given that a consensus on which game theoretic solution concept should be leveraged for a given situation can be elusive, we recommend the consideration of alternative frameworks to ascertain whether identified outcomes are robust to the choice of a solution concept. Should analysis of the prospect game indicate that the predicted persuader behavior is unacceptable, a similar set of regulated prospect games should be solved to provide a collection of intervention options. Such a methodology is illustrated in Section 3.4.

Depending on the nature of influence and regulation, both of the games described herein may involve some form of communication. As such, they are related to signaling (Shoham and Leyton-Brown, 2008), cheap talk (Crawford and Sobel, 1982), persuasion (Milgrom, 1981), and Bayesian persuasion games (Kamenica and Gentzkow, 2011). However, prospect games and regulated prospect games are unique in their utilization of a behavioral theory of choice (i.e., CPT) to describe lower-level decisionmakers’ behaviors. More formally, since many communication games adopt a

leader-follower framework, they can be modeled and solved as bilevel programs. A prospect game has a similar structure albeit with multiple leaders, and its agent-based modeling approach, enabled by CPT, allows for it to be modeled as a standard, single-level, optimization problem, and solved as a normal form game. This incorporation of CPT allows for a similar transformation in regulated prospect games, from a trilevel to a bilevel math programming formulation.

The two modeling constructs described herein incorporate quantitative psychology and decision science theory within a mathematical programming construct to enable the comprehension and regulation of competitive influence with the intent of providing models capable of informing defensive security policy. The remainder of this research details these contributions as follows. Section 3.2 formally defines prospect games; describes how they relate to existing psychological, economic, and neuroeconomic literature; and discusses appropriate solution methodologies. Section 3.3 defines regulated prospect games under a general game theoretic solution concept; provides three mathematical programming formulations of these games under specific game theoretic solution concepts; and explores suitable solution methods under varying baseline assumptions. An illustration of how these models can be used to inform national security policy is provided in Section 3.4, followed by concluding remarks in Section 3.5.

3.2 Prospect Games

Whereas the persuasion programs of Caballero et al. (2018) model a single persuader interacting with a group of decisionmakers, a prospect game (PG) considers multiple persuaders simultaneously influencing the decisionmaker population. More formally, in a PG, multiple persuaders simultaneously compete to alter a decisionmaker(s) decision tree, and/or the decisionmakers' evaluation thereof. The persuaders

act first, and then the decisionmakers (i.e., the populace) select a prospect. Such a structure, as displayed in Figure 8, is consistent with many modern influence campaigns.

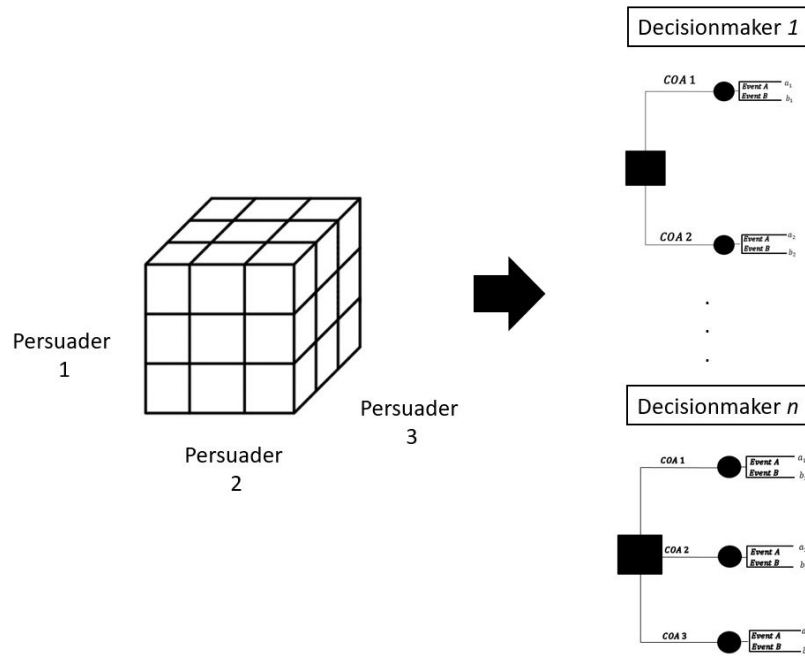


Figure 8. PG with 3 Persuaders each having 3 Actions, and n decisionmakers

Each persuader is self-interested and desires the decisionmakers to each select some respective prospect. These prospects may or may not coincide for any given persuader-and-decisionmaker combination. The objective within a prospect game is not to optimize a single persuader's decision, but to model and understand the strategic interaction between persuaders competing to influence the same populace of decisionmakers. Once this interaction is understood, a decision can be made on whether intervention is necessary. If intervention is required, the techniques described in Section 3.3 describe how a collection of intervention strategies can be assembled.

The leaders-followers framework of a PG naturally lends itself to a bilevel programming formulation. However, when solving a PG, the assumption that the decisionmakers are boundedly rational in accordance with CPT allows for the problem's

transformation to a normal form game with CPT-determined payoffs. The incorporation of such boundedly rational decisionmakers is appropriate, given we model the choices of humans in a descriptive, behavioral sense.

CPT was developed by Tversky and Kahneman (1992) via extensive human-subject testing and has proven itself effective in the representation of human behavior. It is a descriptive theory of choice under risk and uncertainty that is an alternative to the normative Expected Utility Theory. CPT is unique in its findings that humans systematically overweight low probabilities, underweight high probabilities, evaluate utility from a reference point, and experience losses stronger than gains of the same value. Mathematically, these characteristics are represented by an inverse-sigmoid shaped probability weighting function and an asymmetrical sigmoid-shaped utility function (e.g., equations (19a)–(19d)). Although, there exists multiple specific forms of these functions, they are all characterized by some parameters altering the subadditivity of the probability weighting function (e.g., γ_i in equation (19a)), the decisionmaker's loss aversion (e.g., λ_i in equation (19d)) and the utility function curvature (e.g., α_i in equation (19c)). Therefore, CPT provides a flexible, descriptive and deterministic basis for modeling the populace's behavior and, as explained later in this Section 3.2, enables standard game theoretic techniques to be utilized when solving a PG.

In this section, we formally define prospect games by incorporating the following assumptions: (1) persuader action spaces are discrete; (2) the populace members respectively make a decision from among a finite set of actions in accordance with Cumulative Prospect Theory; and (3) decisionmaker prospects, CPT-related parameters, and the effect of exogenous actions on those parameters are known. This final assumption indicates that PGs in the form presented are most useful in situations wherein decisionmaker behavior is well-studied; however, as discussed in Section 3.3,

PGs can be altered to a Bayesian context to accommodate more uncertain environments. Moreover, the models presented herein deliberately maintain generality of the following options to ensure their versatility: (1) type of uncertainty (i.e., risk or ambiguity); (2) form of decisionmakers' CPT utility and probability weighting functions; (3) rationality of persuaders; and (4) specific nature of discrete persuader action.

Defining Prospect Games.

The behavior of persuaders in a PG depends upon the choices of the decisionmakers, which are in turn characterized by CPT. Given the persuaders' actions, CPT provides a deterministic model of the followers' behavior that can be utilized to discern each leader's utility and reduce the game to a normal form setting. Therefore, the modeling of each respective decisionmaker via their CPT parameters and decision trees constitutes the foundation of the game and is accomplished utilizing the following sets, and variables.

Sets

\mathcal{I} : Decisionmakers upon which persuasion is being conducted of the form

$$\mathcal{I} = \{1, 2, \dots\}$$

\mathcal{P} : Set of persuaders of the form $\mathcal{P} = \{1, 2, \dots\}$

\mathcal{A} : Set of persuader action profiles such that $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \dots$, where \mathcal{A}_p is a finite set of actions available to persuader p

\mathcal{S} : Set of persuader mixed-strategy profiles such that $\mathcal{S} = \mathcal{S}_1 \times \mathcal{S}_2 \times \dots$, where \mathcal{S}_p is the set of all probability distributions over \mathcal{A}_p

\mathcal{J}_i : Set of n_i prospects offered to each member i of the form $\mathcal{J}_i = \{1, \dots, n_i\}$

\mathcal{K}_{ij} : Set of m_{ij} events (i.e., outcomes) for member i in prospect j of the form

$$\mathcal{K}_{ij} = \{1, \dots, m_{ij}\}$$

Ω_u : Some subset of \mathcal{K}_{ij} such that $\{\Omega_u \in \mathcal{K}_{ij} : 1 \leq |\Omega_u| \leq m_{ij}\}$

CPT Variables

x_{ijk} : Gain/loss for member i for k^{th} event of prospect j after $\mathbf{a} \in \mathcal{A}$

p_{ijk} : Probability of k^{th} event of prospect j for member i after $\mathbf{a} \in \mathcal{A}$;
used in conditions of risk

$p_{ij\Omega_u}$: Judged probability that the disjunction of events in Ω_u occurs after
 $\mathbf{a} \in \mathcal{A}$; *used in conditions of ambiguity*

γ_i : Gain probability weighting coefficient for member i after $\mathbf{a} \in \mathcal{A}$

δ_i : Loss probability weighting coefficient for member i after $\mathbf{a} \in \mathcal{A}$

α_i : Gain utility curvature coefficient for member i after $\mathbf{a} \in \mathcal{A}$

β_i : Loss utility curvature coefficient for member i after $\mathbf{a} \in \mathcal{A}$

λ_i : Loss aversion coefficient for member i after $\mathbf{a} \in \mathcal{A}$

r_i : Reference point for member i after $\mathbf{a} \in \mathcal{A}$

A given collective action $\mathbf{a} \in \mathcal{A}$ by the persuaders establishes the decision setting for each decisionmaker. We allow for the underlying probability weighting, utility curvature, loss aversion, and reference point parameters to be altered, a modeling choice we discuss in Section 3.2. Henceforth, the caret notation distinguishes between baseline CPT parameters and their updated values after action \mathbf{a} (e.g., $\hat{\gamma}_i$ and

γ_i , respectively). Moreover, for each prospect $j \in J_i$, persuasive actions may alter the outcome value (i.e., from \hat{x}_{ijk} to x_{ijk}) or the uncertainty (i.e., from \hat{p}_{ijk} to p_{ijk} and from $\hat{p}_{ij\Omega_u}$ to $p_{ij\Omega_u}$) associated with any $k \in \mathcal{K}_{ij}$. The variable p_{ijk} is utilized under conditions of risk, whereas $p_{ij\Omega_u}$ is required for conditions of ambiguity in accordance with the belief-based model of CPT by Fox and Tversky (1998). The generic structure of a prospect game allows for each of these parameters to be updated as a function of \mathbf{a} ; however, some may be assumed to be constant without loss of generality. The psychological and neuroeconomic justification for such update functions is also discussed in Section 3.2.

Each decisionmaker discerns the value for every prospect in accordance with CPT as indicated in equations (11)–(17) (Tversky and Kahneman, 1992). A reader familiar with decision science may recognize equation (11) as the sum of two Rank-dependent Utility (RDU) functionals. CPT can be viewed as a generalization of RDU that incorporates all three components of risk attitude: utility curvature, probabilistic sensitivity, and loss aversion (Wakker, 2010).

$$V_{ij}(\mathbf{a}) = V_{ij}^+(\mathbf{a}) + V_{ij}^-(\mathbf{a}), \text{ where} \quad (11)$$

$$V_{ij}^+(\mathbf{a}) = \sum_{l=1}^z \pi_{ijl}^+(\mathbf{a}) v^+(x_{ijl}, \mathbf{a}), \quad (12)$$

$$V_{ij}^-(\mathbf{a}) = \sum_{l=-y}^0 \pi_{ijl}^-(\mathbf{a}) v^-(x_{ijl}, \mathbf{a}), \quad (13)$$

$$\pi_{ijz}^+(\mathbf{a}) = W^+(E_{ijz}, \mathbf{a}), \quad (14)$$

$$\pi_{ij(-y)}^-(\mathbf{a}) = W^+(E_{ij[-y]}, \mathbf{a}), \quad (15)$$

$$\pi_{ijl}^+(\mathbf{a}) = W^+(E_{ijl} \cup \dots \cup E_{ijz}, \mathbf{a}) - W^+(E_{ij[l+1]} \cup \dots \cup E_{ijz}, \mathbf{a}), \quad 0 \leq l \leq z - 1 \quad (16)$$

$$\pi_{ijl}^-(\mathbf{a}) = W^+(E_{ij[-y]} \cup \dots \cup E_{ijl}, \mathbf{a}) - W^+(E_{ij(-y)} \cup \dots \cup E_{ij[l-1]}, \mathbf{a}), \quad -y \leq l \leq 0. \quad (17)$$

For each individual i , a prospect j , having m_{ij} total outcomes, is separated into z gain outcomes and y loss outcomes, each having an associated value x_{ijl} gain/loss from the reference point. The value of the prospect, $V_{ij}(\mathbf{a})$, is calculated as the sum of the gain and loss values (e.g., $V_{ij}^+(\mathbf{a})$ and $V_{ij}^-(\mathbf{a})$) that are respectively calculated as the sumproduct of π_{ijl}^\pm and v^\pm . That is, each outcome's gain and loss values are determined via the decision weight and utility functions (e.g., equations (12) and (13)). The decision weights are calculated as the marginal difference of the event weighting function, W^+ or W^- , for gain and loss ranks, respectively, wherein a gain rank of event k is the probability of receiving a better outcome, and a loss rank is the probability of a worse outcome (Wakker, 2010). Likewise, the value function is generally modeled in a piecewise nature allowing for different functional forms of gains and losses (i.e., v^+ and v^-).

Based on the results of Fox and Tversky (1998), the probability weighting functions, w^\pm , can be used as the event weighting function, W^\pm , in conditions of risk or uncertainty. However, these results hold under the assumption that judged probabilities satisfy a non-extensional theory of subjective probability (i.e., support theory). Therefore, the specific probabilities utilized in the weighting functions under conditions of risk are different than those under uncertainty, as seen in equations (18a) and (18b), respectively. The difference between the two equations derives from the preservation of the extensionality property under risk and its loss under uncertainty. That is, under uncertainty, the subjective probability of event $E_{ijl} \cup \dots \cup E_{ijs}$ must be solicited directly whereas, under conditions of risk, it can be calculated as the sum of p_{ijl}, \dots, p_{ijs} .

$$W^\pm(E_{ijl} \cup \dots \cup E_{ijs}, \mathbf{a}) = w^\pm\left(\sum_{t=l}^s p_{ijt}, \mathbf{a}\right) \quad (18a)$$

$$W^\pm(E_{ijl} \cup \dots \cup E_{ijs}, \mathbf{a}) = w^\pm\left(p_{ij[E_{ijl} \cup \dots \cup E_{ijs}]}, \mathbf{a}\right) \quad (18b)$$

Maintaining consistency with Caballero et al. (2018), we adopt the original functions of Tversky and Kahneman (1992) in our worked examples, as depicted in equations (19a)–(19d). However, they are not fundamental to the modeling approach examined herein, and an interested researcher may chose to adopt alternatives forms for either the utility or probability weighting functions (e.g., see Prelec, 1998; Wakker, 2010)

$$w^+(p_{ijk}, \mathbf{a}) = \frac{p_{ijk}^\gamma}{(p_{ijk}^\gamma + (1 - p_{ijk'})^\gamma)^{\gamma-1}}, \quad (19a)$$

$$w^-(p_{ijk}, \mathbf{a}) = \frac{p_{ijk}^\delta}{(p_{ijk}^\delta + (1 - p_{ijk'})^\delta)^{\delta-1}}, \quad (19b)$$

$$v^+(x_{ijk}, \mathbf{a}) = x_{ijk}^{\alpha_i}, \quad (19c)$$

$$v^-(x_{ijk}, \mathbf{a}) = -\lambda(-x_{ijk})^{\beta_i}. \quad (19d)$$

Each player i is assumed to select the prospect j that maximizes equation (11). The collective selection of prospects by all players in \mathcal{I} determines the persuaders' payoffs. This information is determined by the variables $\phi_{ip} = f(V_i(\mathbf{a}), q_{ip}), \forall i \in \mathcal{I}, p \in \mathcal{P}$ where $V_i(\mathbf{a})$ is the CPT evaluation of all prospects J_i , q_{ip} is the prospect persuader p desires decisionmaker i to select, and $f : V_i(\mathbf{a}) \times q_i \rightarrow \{0, 1\}$. It is helpful to envision ϕ_{ip} as constituting elements in a matrix Φ wherein a value of one in the entry (i, p) indicates that decisionmaker i prefers the prospect desired by persuader p , and zero otherwise.

For some $a \in \mathcal{A}$, the utility vector for the persuaders is defined as

$$\theta_{\mathbf{a}} = \{g_1(\Phi, \mathbf{a}), g_2(\Phi, \mathbf{a}), \dots, g_{|\mathcal{P}|}(\Phi, \mathbf{a})\}, \quad (20)$$

such that $g_p(\Phi, \mathbf{a})$ calculates a persuader's utility for the decisionmaker preferences in

Φ and some collective action $\mathbf{a} = (a_1, a_2, \dots)$ wherein each persuader p independently controls the variable a_p . A given persuader's utility can be determined based upon a myriad of factors, such as whether a decisionmaker prefers prospect q_{ip} or acted in accordance to a competitor's desire, as well as the amount of resources expended in influential action compared to other persuaders. The $\theta_{\mathbf{a}}$ -values can be aggregated to form the set of vectors $\Theta = \{\theta_{\mathbf{a}}, \forall \mathbf{a} \in \mathcal{A}\}$. Therefore, since the action space \mathcal{A} is finite, we can define a normal form game via the tuple $(\mathcal{P}, \mathcal{A}, \Theta)$.

On the Malleability of Prospects and CPT Parameters.

Human decisions are shaped by subjective evaluations of uncertainty and value, emotions, and a variety of other psychological phenomena (Kahneman, 2011; Lerner et al., 2015). Therefore, persuader actions in this research are allowed to affect the perception of an outcome's likelihood, the perception of an outcome's value, and/or an individual's CPT-parameter values, and in this way the persuaders are able to affect the populace's evaluation of their respective decision trees. Because such a perspective does not have consensus among decision-science scholars, we briefly review the neuroscientific literature supporting this point of view. However, for a more thorough literature review spanning multiple disciplines (e.g., neuroeconomics, psychology, and management science), we refer the reader to Caballero et al. (2018).

The empirical evidence from economics studies suggesting the dynamic nature of risk attitude is reinforced by pharmacological experiments in neuroscience. Crockett and Fehr (2014) provided an overview of the evidence suggesting that the release of certain neuromodulators affect behavior in risky choice. Studies have shown dopamine to promote risky-choice (Imamura et al., 2006; Eisenegger et al., 2010; Crockett and Fehr, 2014), and the ingestion of amphetamines has been shown to affect behavior analogously (Onge and Floresco, 2009). Other research efforts are ongoing to deter-

mine how different neuromodulators, such as serotonin and norepinephrine, contribute to establishing risk attitudes (Crockett and Fehr, 2014).

Louie and De Martino (2014) provided support to the malleable nature of risky choice by invoking the neurobiological *efficient coding hypothesis* (Barlow, 1961). This hypothesis postulates that, due to the brain's inherent biophysical constraints (e.g., neuron firing rate and metabolic costs), efficient neural input-output functions are formed based on some expected range of input values. Should the range of input parameters change, the associated neural function changes as well, in a process falling under the general category of *gain control* (Louie and De Martino, 2014). This underlying behavior of the brain at the neuronal level is related to the results of Stewart et al. (2006) and Stewart et al. (2015), who linked phenomena explained by CPT with an individual performing comparisons to previous experiences in working memory.

Taken collectively, these studies suggest some malleable, neurological phenomena drives risk attitudes, and the impact of a persuasive action on it can be estimated. Therefore, we proceed with the assumption that these factors have been estimated via human subject testing; it is our intent that the following models and accompanying solutions methods presented demonstrate the utility that can be gained from such human subject testing and thereby motivate it within the discipline.

Solving Prospect Games.

The solution concept utilized for the game $(\mathcal{P}, \mathcal{A}, \Theta)$ depends upon the persuaders' assumed rationality. If they are assumed to be perfectly rational, then traditional game theoretic concepts are applicable (e.g., Nash equilibrium, ε -Nash equilibrium, trembling hand equilibrium, maxmin or minmax payoffs, et cetera), depending on the context (Shoham and Leyton-Brown, 2008). However, if the persuaders are assumed to be boundedly rational, solution concepts from behavioral game theory, such as

the cognitive hierarchy model (Camerer et al., 2004) or quantal response equilibrium (McKelvey and Palfrey, 1995), are appropriate and should be compared to solutions found via perfect rationality analysis. In either scenario, these solution concepts are expressed in terms of mixed strategies. That is, the solution is represented as a collection of probability distributions $S \in \mathcal{S}$. In a given solution, each persuader's strategy is probabilistic and denoted as the vector S_p composed of the elements S_{pl} for each $l \in \mathcal{A}_p$.

In addition to the difference in underlying rationality assumptions, the algorithms utilized in a perfectly rational or a boundedly rational setting have distinct characteristics. For example, the computation of a Nash equilibrium for two-player games can be computed efficiently as a linear program or a linear complementarity problem in the zero-sum and general-sum cases, respectively. However, an n -player, general-sum game requires a more complex formulation as a nonlinear complementarity problem. Games also may have multiple Nash equilibrium and, in the worst-case, even a two-player, general-sum game requires a solution time that is exponential in the number of actions of each player (Shoham and Leyton-Brown, 2008). Moreover, even if all of these equilibriums are found, it is difficult to distinguish among them systematically. From this standpoint, the maxmin or minmax payoffs solution concepts are beneficial in that they can be computed quickly and provide a single prediction (Shoham and Leyton-Brown, 2008). Such characteristics are shared by many boundedly rational solution concepts. For instance, the cognitive hierarchy model (Camerer et al., 2004) can be computed efficiently and yields a single prediction for a given set of assumptions.

In practice, it may be advisable to examine many games under varying assumptions of persuader effects. The utility garnered, by both persuader and decisionmaker, in some competitive persuasion scenarios is difficult to discern and must be estimated.

Therefore, inputs utilized by any solution concept involve an accumulation of experimental error. By analyzing a variety of games that constitute some confidence interval of utility and persuasion effects, a regulator is able to generate an expected range of behavior and ensure greater robustness with respect to estimation error.

Finally, multiple solution concepts should be considered when analyzing a prospect game, especially when the goal is to inform a regulator of expected outcomes under null action. Generally speaking, multiple solution concepts should be examined to garner collective insights since no single solution method dominates all others but instead depend upon specific rationality assumptions. Therefore, in Section 3.4, we utilize the solution concepts of the Nash equilibrium, correlated equilibrium, and cognitive hierarchy model in an effort to understand persuader interaction.

3.3 Regulated Prospect Games

Regulated prospect games (RPGs) are an extension of PGs wherein an external regulator attempts to affect persuader behavior in equilibrium. Figure 9 illustrates this framework.

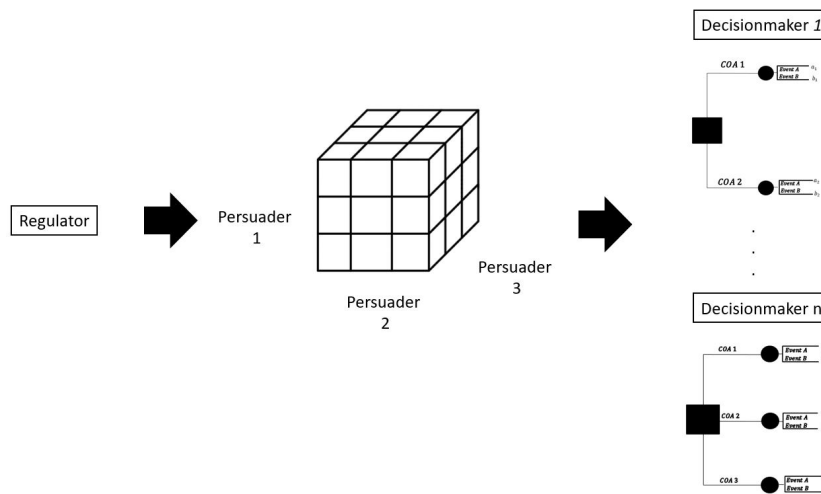


Figure 9. Structure of Regulated Prospect Games

We allow for the strategy space of a regulator to be continuous or discrete. As

with the persuaders, no assumptions are made regarding specific means by which their actions modify decisionmaker or persuader behavior. Such actions may change the structure of the game or, potentially, the psychology of subsequent agents when selecting a course of action.

An RPG can be considered a trilevel program wherein the regulator, persuaders and the decisionmakers are the first, second and third level entities, respectively. However, in this section, we describe how the PG reduction from the previous section can be utilized to transform an RPG to a bilevel program and how the utilization of specific game theoretic solution concepts enables further transformation to a single-level form.

Examining the special cases of RPGs under the same three solution concepts (i.e., Nash equilibria, correlated equilibria, and cognitive hierarchy), we examine how the problems can be solved. For some instances, a commercial solver is readily able to solve an RPG. However, for those instances wherein a commercial solver proves insufficient, heuristic methods to attain quality, feasible solutions for the regulator are discussed.

Defining Regulated Prospect Games.

A regulator must choose some action \mathbf{b} from a set of possible intervention strategies, \mathcal{B} , to affect the game $(\mathcal{P}, \mathcal{A}, \Theta)$. The regulator may impact actions available to the persuaders, alter persuader payoffs directly or via decisionmaker preferences effects, or in the case of boundedly rational persuaders, their ability to think strategically.

An RPG maintains the same payoff structure as a PG for both decisionmakers and persuaders, but it adds an additional element pertaining to the regulator's payoff. RPGs levy an additional layer of complexity to PGs, as their task is, essentially, to modify the game to be played by the persuaders. Therefore, regardless of the assumed

rationality of the persuaders, an RPG can be formulated as a bilevel program in the following manner:

$$\begin{aligned}
& \max_{\mathbf{b}} h_1(S, \mathbf{b}) \\
& \text{subject to } \mathbf{b} \in \mathcal{B}, \\
& \max_S \sum_{\mathbf{a} \in \mathcal{A}} g_1(\Phi, \mathbf{a}, \mathbf{b}) \prod_{p \in \mathcal{P}} S_{p\mathbf{a}_p}; \sum_{\mathbf{a} \in \mathcal{A}} g_2(\Phi, \mathbf{a}, \mathbf{b}) \prod_{p \in \mathcal{P}} S_{p\mathbf{a}_p}; \dots \\
& \text{s.t.: } S \in \mathcal{S}(\mathbf{b}),
\end{aligned}$$

where

$$\phi_{ip} = f(V_i(\mathbf{a}, \mathbf{b}), q_{ip}), \forall i \in \mathcal{I}, p \in \mathcal{P}. \quad (21)$$

The upper level problem dictates the regulator take an action $\mathbf{b} \in \mathcal{B}$ to maximize the objective function $h_1(S, \mathbf{b})$. In turn, this action can affect the set of feasible persuader strategies, $\mathcal{S}(\mathbf{b})$. For our purposes, we assume all possible actions of a player p are included in \mathcal{A}_p , and a regulator action \mathbf{b} can only remove, not add, from this set. Regulator action, in conjunction with the persuaders' actions, may also alter the decisionmaker preferences, and this effect is modeled generically by representing the CPT prospect valuation vectors as $V_i(\mathbf{a}, \mathbf{b})$. Likewise, to account for instances wherein the regulating action has a direct impact on a persuaders payoff (e.g., monetary fines) the vector of regulating actions, \mathbf{b} , is represented in the persuaders' payoff functions as $g_p(\Phi, \mathbf{a}, \mathbf{b})$. For a given $\mathbf{b} \in \mathcal{B}$, the persuaders' payoffs are determined as in a PG. Thus, the lower-level optimization problem is decentralized; each persuader p controls a subset of the variables, S_p , and maximizes their respective objectives

simultaneously.

An RPG has a PG embedded in its constraints, and can be viewed as a mechanism design problem with a singleton set of joint type vectors. However, should a regulator decide to analyze a set of games, such as those formed from a range of estimates with respect to utility and persuasion effects, and calculate a probability distribution over this set, the cardinality of the joint type vector can be expanded. In this case, the RPG exists in a Bayesian setting and is more reminiscent of traditional mechanism design problems. The resulting math programs are very similar except for a few additional constraints. That is, the inequalities illustrated in the following programs must be repeated for every game such that each has its own corresponding block of constraints, and the objective function must be modified to an expected value computation.

The bilevel program representing an RPG can be analyzed in a variety of ways depending upon both persuader rationality assumptions and the adopted solution concept. We continue by illustrating three special cases corresponding to persuaders behaving in accordance with the Nash equilibrium, correlated equilibrium, and cognitive hierarchy model solution concepts.

Perfect Rationality: Nash Equilibrium.

If it is assumed the persuaders are perfectly rational, and it is desired to discern regulator action under a Nash equilibrium, the RPG can be reformulated as a single-level program by adapting the framework of Shoham and Leyton-Brown (2008) with regard to n -player, general sum games.

$$\max_{S, \mathbf{b}} h_1(S, \mathbf{b}) \quad (22a)$$

$$\text{subject to } \mathbf{b} \in \mathcal{B}, \quad (22b)$$

$$\sum_{\mathbf{a} \in \mathcal{A} | a_p = l} g_p(\Phi, \mathbf{a}, \mathbf{b}) \prod_{z \neq p} S_{z a_z} - \sum_{\mathbf{a} \in \mathcal{A}} g_p(\Phi, \mathbf{a}, \mathbf{b}) \prod_{z \in P} S_{z a_z} \leq 0, \quad (22c)$$

$$\forall p \in \mathcal{P}, l \in \mathcal{A}_p,$$

$$\sum_{l \in \mathcal{A}_p} S_{pl} = 1, \quad \forall p \in P, \quad (22d)$$

$$S_{pl} \leq h_2(\mathbf{b}, S_{pl}), \quad \forall p \in P, l \in \mathcal{A}_p, \quad (22e)$$

$$S_{pl} \geq 0, \quad \forall p \in P, l \in \mathcal{A}_p, \quad (22f)$$

where

$$\phi_{ip} = f(V_i(\mathbf{a}, \mathbf{b}), q_{ip}), \quad \forall i \in \mathcal{I}, p \in \mathcal{P}.$$

In this reformulation, the left-hand side of constraint (22c) represents the effect of persuader p deviating from their respective S_p strategy to the pure strategy l . If it is negative-valued, the persuader is better off not deviating. Therefore, if constraint (22c) is satisfied for all persuaders and available actions, then no one has incentive to deviate from their respective strategy S_p , and S is a Nash equilibrium. The constraints (22d)–(22f) ensure a persuader's mixed strategy profile abides by the tenets of probability theory, and they also allow for the possibility of a regulator removing some action l in \mathcal{A}_p for persuader p . That is, $h_2(\mathbf{b}, S_{pl})$ is assumed to utilize the regulator action (i.e., \mathbf{b}) and an element of a persuader's mixed strategy (i.e., S_{pl}) as inputs, and it maps them to the binary outputs, $\{0,1\}$, that form an upper bound on S_{pl} . Therefore, if this mapping yields an output of zero, the regulator has removed

action l from consideration by persuader p .

Furthermore, this formulation implicitly adopts an *optimistic* regulator as detailed in the bilevel programming literature (e.g., see Dempe, 2002; Colson et al., 2007), and it assumes the persuaders will adopt the Nash equilibrium that is most beneficial to the upper level objective function. All other Nash equilibriums are feasible, but this optimistic assumption allows the regulator to distinguish between them, and their associated objective function values. By setting the \mathbf{b} -values that yield this optimistic solution as parameters and resolving the program by minimizing the objective function instead of maximizing it, the regulator is able to determine the associated solution's *pessimistic* solution. In this manner, the regulator is able to identify the down-side risk associated with their optimistic solution.

Perfect Rationality: Correlated Equilibrium.

If persuaders are perfectly rational, but it is desired to discern regulator action under the correlated equilibrium solution concept, the bilevel program can be simplified further. A correlated equilibrium can be interpreted as relating to a public signal. That is, if all players in a game prefer to abide by the “recommendations” provided by an external agent, the resulting profile is a correlated equilibrium. In this setting, the regulator's task is to determine a probability vector S' composed of elements S'_a for every $\mathbf{a} \in \mathcal{A}$ such that no player prefers to deviate from \mathbf{a} when it is signaled. Utilizing the notation $\mu^{\mathbf{a}}$ to signify the action profile \mathbf{a} with a unilateral deviation of $\mu_p^{\mathbf{a}}$ by player p , and the function $h_3(\mathbf{a}, \mathbf{b}) \rightarrow \{0, 1\}$ to indicate whether regulator action \mathbf{b} removed the possibility of persuader action \mathbf{a} , the problem can be expressed as follows the following program (Shoham and Leyton-Brown, 2008):

$$\max_{S', \mathbf{b}} h_1(S', \mathbf{b}) \quad (23a)$$

$$\text{subject to } \mathbf{b} \in \mathcal{B}, \quad (23b)$$

$$\sum_{\mathbf{a} \in \mathcal{A} | a_p = l} S'_a g_p(\Phi, \mathbf{a}, \mathbf{b}) \geq \sum_{\mathbf{a} \in \mathcal{A} | a_p = l} S'_a g_p(\Phi, \mu^{\mathbf{a}}, \mathbf{b}) h_3(\mu^{\mathbf{a}}, \mathbf{b}), \quad (23c)$$

$$\forall p \in \mathcal{P}, l \in \mathcal{A}_p, \mu_p^{\mathbf{a}} \neq l$$

$$S'_a \leq h_3(\mathbf{a}, \mathbf{b}), \quad \forall a \in \mathcal{A} \quad (23d)$$

$$S'_a \geq 0, \quad \forall a \in \mathcal{A}, \quad (23e)$$

$$\sum_{a \in \mathcal{A}} S'_a = 1. \quad (23f)$$

where

$$\phi_{ip} = f(V_i(\mathbf{a}, \mathbf{b}), q_{ip}), \quad \forall i \in \mathcal{I}, p \in \mathcal{P}.$$

Constraint (23b) ensures the decision variables S' and \mathbf{b} constitute a correlated equilibrium. Constraints (23d)–(23f) constrain the signaled probabilities analogously to how constraints (22d)–(22f) bound the persuaders' mixed strategies. However, due to the signaling nature of a correlated equilibrium, the assumption of optimism is no longer required as in the previous Nash equilibrium formulation. Although a game may have many correlated equilibria, the probability vector S' is actually controlled by the regulator and they are free to select it in accordance with their preference.

Bounded Rationality: Cognitive Hierarchy.

The math programming reformulations presented thus far in this section focus on the case of perfectly rational persuaders. Such a tendency for optimization models to

assume perfectly rationality is due to the fact that many boundedly rational solution concepts are heuristic in nature; they are designed to replicate imperfect psychological processes that do not always translate, in a strict sense, to an optimization framework. However, in this section, a boundedly rational solution concept is considered that maintains the concept of utility maximization but relaxes assumptions regarding an individual's understanding of his opponents that allows for modeling via mathematical programming formulations.

The cognitive hierarchy model (Camerer et al., 2004) assumes players are defined by the number of reasoning steps they compute in selecting an action, and their beliefs with regard to the number of steps utilized by their opponents. Players do not believe other agents can use as many reasoning steps as they do (i.e., a κ -step player believes their opponents are $[\kappa-1]$ -step players or less). A game is assumed to be characterized by a true probability distribution over the number of reasoning steps a player utilizes. Generally, a Poisson distribution defined by the parameter τ is used. This true distribution, $\hat{f}(\zeta)$, is assumed to be perceived accurately by a κ -step player but normalized to inform their beliefs about their opponents (i.e., their beliefs of the proportion of h -step players are represented as $\rho_{\kappa h} = \hat{f}^{(h)} / \sum_{\zeta < \kappa} \hat{f}(\zeta)$).

Therefore, a structure similar to that utilized for modeling persuader behavior under the Nash equilibrium solution concept can be adopted as follows with the incorporation of a set of new decision variables, S_{pl}^{κ} , that indicate the probability of persuader p employing strategy l when utilizing κ -steps of thought. This model assumes play will occur in accordance with the cognitive hierarchy model's expectation over a population defined by τ .

$$\max_{S, \mathbf{b}} h_1(S, \mathbf{b}) \quad (24a)$$

$$\text{subject to } \mathbf{b} \in \mathcal{B}, \quad (24b)$$

$$\tau = h_4(\mathbf{b}, \hat{\tau}) \quad (24c)$$

$$\rho_{\kappa h} = \frac{\hat{f}(\tau, h)}{\sum_{y < \kappa} \hat{f}(\tau, y)}, \quad \forall \kappa < M, h < \kappa, \quad (24d)$$

$$S_{p\mathbf{a}_p}^0 = \frac{1}{|\mathcal{A}_p|}, \quad \forall p \in \mathcal{P}, \quad (24e)$$

$$\sum_{h < \kappa} \sum_{\mathbf{a} \in \mathcal{A}_{|a_p=l}} \rho_{\kappa h} g_p(\Phi, \mathbf{a}, \mathbf{b}) \prod_{z \neq p} S_{z\mathbf{a}_z}^h - \sum_{h < \kappa} \sum_{\mathbf{a} \in \mathcal{A}} \rho_{\kappa h} g_p(\Phi, \mathbf{a}, \mathbf{b}) \prod_{z \in \mathcal{P}} S_{z\mathbf{a}_z}^h \leq 0,$$

$$\forall p \in \mathcal{P}, l \in \mathcal{A}_p, \kappa < M, \quad (24f)$$

$$S_{pl} = \sum_{\kappa < M} \hat{f}(\kappa) S_{pl}^\kappa, \quad \forall p \in \mathcal{P}, l \in \mathcal{A}_p, \quad (24g)$$

$$\sum_{l \in \mathcal{A}_p} S_{pl}^\kappa = 1, \quad \forall p \in \mathcal{P}, \kappa < M, \quad (24h)$$

$$S_{pl} \leq h_2(\mathbf{b}, S_{pl}), \quad \forall p \in \mathcal{P}, l \in \mathcal{A}_p, \quad (24i)$$

$$S_{pl}^\kappa \geq 0, \quad \forall p \in \mathcal{P}, l \in \mathcal{A}_p, \kappa < M \quad (24j)$$

$$S_{pl} \geq 0, \quad \forall p \in \mathcal{P}, l \in \mathcal{A}_p, \quad (24k)$$

where

$$\phi_{ip} = f(V_i(\mathbf{a}, \mathbf{b}), q_{ip}), \quad \forall i \in \mathcal{I}, p \in \mathcal{P}.$$

It has been shown that different game structures can generate alternative values of τ (Camerer et al., 2004). This behavior is modeled via constraint (24c) by assuming the function $h_4(\mathbf{b}, \hat{\tau})$ updates this value based on regulator action, and the value $\hat{\tau}$ associated with the baseline persuasion game (i.e., null regulator action). This

value is utilized in constraint (24d) to update the perceived probabilities of a κ -level player wherein $\hat{f}(\tau, h)$ is the Poisson probability density function of h with rate τ . Constraint (24e) ensures all 0-step players randomize uniformly and constraint (24f) ensures a κ -step persuader selects their perceived best response. The strategies of all κ -step players are then aggregated as in Camerer et al. (2004) to provide the behavioral predictions of play in the cognitive hierarchy model via constraint (24g). The remaining constraints ensure the tenets of probability are not violated and allow for the regulator to remove strategies via the function $h_2(\mathbf{b}, S_{pl})$ as defined previously.

Camerer et al. (2004) assumed for simplicity that players would “randomize equally if two or more strategies have identical expected payoffs” but acknowledged the cognitive hierarchy model could address other methods. Therefore, no explicit assumptions are made herein regarding how players distinguish between strategies yielding equal payoffs.

Solving Regulated Prospect Games.

Optimally solving one of the aforementioned mathematical programs is not a simple task. To illustrate this point, first consider the determination of ϕ_{ip} via equation (21). The function f maps the persuaders’ action vector, \mathbf{a} , and the regulator’s action vector, \mathbf{b} , into a binary response for a given persuader p and a decisionmaker i . Therefore, aside from the trivial case wherein ϕ_{ip} is constant, the function will exhibit discontinuities when ϕ_{ip} switches to either 0 or 1. The functions $g_p(\Phi, \mathbf{a}, \mathbf{b})$ will exhibit similar discontinuities because they depend on these ϕ_{ip} variables. Moreover, the function f implicitly concerns a maximization operation coupled with a decision problem. That is, for some decisionmaker i and persuader p if $q_{ip} = \operatorname{argmax}_j \{V_{ij}(\mathbf{a}, \mathbf{b})\}$, then $\phi_{ip} = 1$ and zero otherwise.

These difficulties are compounded by the nonlinearities associated with determin-

ing the values in the vector $V_i(\mathbf{a}, \mathbf{b})$. To calculate these values, a sorting operation that orders the prospects according to ascending outcome value is required for use in equations (14)–(17). Once accomplished, further nonlinearities are encountered in equations (19a)–(19d), albeit via an effect that depends on what CPT parameters are assumed to be malleable by the persuaders and regulator.

Even if these difficulties are set aside, the underlying problem of ascertaining a game theoretic solution (e.g., a Nash equilibrium) may be a non-trivial endeavor (Shoham and Leyton-Brown, 2008). Therefore, solving the induced PG for a given \mathbf{b} may itself be a difficult problem. This property is intensified in an RPG, as the regulator is trying to build an optimal game with these equilibrium conditions in the constraints.

Therefore, even though the characterization is ultimately dependent on the form of the regulator action space and the possible effects of regulator or persuader action on decisionmakers, in many RPG instances a global commercial solver is required to find an optimal solution due to the aforementioned difficulties. However, in some instances, even these tools may be inadequate.

Let us first consider a situation for which a global commercial solver can be applied effectively. If a regulator is unable to affect decisionmakers' preferences, the payoff structure of the underlying PG can be calculated by determining the effect of each $\mathbf{a} \in \mathcal{A}$ on Φ . With this the baseline game structure, Φ is fixed and the agents' payoffs (regulator or persuader) depend only on \mathbf{b} and S (or S'). The three mathematical programs previously introduced can then be solved to determine the optimal punitive action \mathbf{b} and the associated game theoretic solution value.

Conversely, if the regulator is able to affect decisionmaker preferences, the instance may be more difficult to solve optimally. To illustrate this property, consider that each outcome $k \in \mathcal{K}_{ij}$ for every decisionmaker i and prospect j , must be sorted

and classified as a gain or loss in order to determine each ϕ_{ip} . Caballero et al. (2018) constructed a set of constraints utilizing matrices of binary decision variables that, when satisfied, ensured the CPT-valuations are calculated correctly. However, commercial solvers are not designed to sort values and generally perform inefficiently at the task. The technique utilized in Caballero et al. (2018) required $2|\mathcal{K}_{ij}|^2$ binary indicator variables for each combination of $i \in \mathcal{I}$ and $j \in \mathcal{J}_i$, and their technique becomes computationally burdensome as instances grow large. Therefore, alternative methodologies are proposed to solve these RPGs, depending upon the form of the regulator's action space.

If the regulator action is discrete, we can avoid the computational difficulty associated with determining ϕ_{ip} by enumerating all possible games a regulator action can create. By calculating off line the decisionmakers' choices for each $\mathbf{a} \in \mathcal{A}$ and $\mathbf{b} \in \mathcal{B}$, and replacing $g_p(\Phi, \mathbf{a}, \mathbf{b})$ with $\lambda_p(\Phi, \mathbf{a}, \nu^{\mathbf{b}})$ as seen in equation (25), the computational difficulties associated with CPT are not encountered by the solver. Instead, their dynamics are incorporated in the off line calculations that determine persuader utility. Formally, the three special case RPGs can be modified by introducing

$$\lambda_p(\Phi, \mathbf{a}, \nu^{\mathbf{b}}) = \sum_{b \in \mathcal{B}} \nu_b g_p(\Phi, \mathbf{a}, \mathbf{b}), \quad (25)$$

and adding the constraints

$$\sum_{b \in \mathcal{B}} \nu_b = 1; \quad \nu_b \in \{0, 1\} \quad \forall \mathbf{b} \in \mathcal{B}.$$

Via off line calculation, the ϕ_{ip} values are determined for each b , and the corresponding $g_p(\Phi, \mathbf{a}, \mathbf{b})$ are calculated for each persuader b and action profile \mathbf{a} . These values can then be input into the RPG formulations as parameters and constraints (22b), (23b), or (24b) can be removed, as appropriate, because the regulator's action space has been enumerated.

If regulator action is continuous, an RPG no longer pertains to the selection of a game among a countable set, but rather optimization among a continuum of games. Therefore, the off line calculation method can no longer be leveraged as a simplification tool. The dynamics of CPT-value calculation would need to be computed within the solver and, as previously discussed, such a procedure is inefficient, especially when there exists a large group of decisionmakers. Therefore, to solve such RPGs, the action space can be approximated via discretization, or an alternative heuristic can be utilized.

A variety of heuristic or meta-heuristic methods can be utilized to solve an RPG (e.g., genetic algorithm or simulated annealing); however, since the objective function depends upon the equilibrium (or alternative game theory solution concept) profile, these methods still must solve a PG for each candidate solution. In light of the aforementioned computational limitations in solving equilibrium problems, we suggest \mathcal{B} be searched via efficient sampling, and advocate for the use of a direct search algorithm such as the Generalized Pattern Search (Lewis and Torczon, 1999) or Mesh Adaptive Direct Search (Audet and Dennis Jr., 2006).

3.4 Example Application: Deterring Electoral Interference

In this section, we consider a notional example that illustrates the utility of PGs and RPGs, and we demonstrate the analysis and insights they can generate for the regulation of competitive persuasive interactions. More specifically, we consider the problem of a democratic government deterring improper interference in its presidential election. Such an application is particularly appropriate in a complete information environment given the large amount of human-subject testing conducted with regard to voting behavior.

The regulator is a democratic government; the persuaders are three organizations

having vested interests in the electoral outcome, identified as Entities 1-3; and the decisionmakers are a subset of targeted voters, as depicted in Figure 10. In this scenario, we assume the regulator exerts limited (i.e., not authoritarian) control and, as such, the government is unable to completely disallow any persuader strategy (i.e., $h_2(\mathbf{b}) = h_3(\mathbf{b}) = 1, \forall \mathbf{b} \in \mathcal{B}$). Instead, the persuaders must be incentivized to take a desired course of action. Likewise, given the political nature of an election, the regulator is assumed to take no action that affects decisionmaker preferences directly.

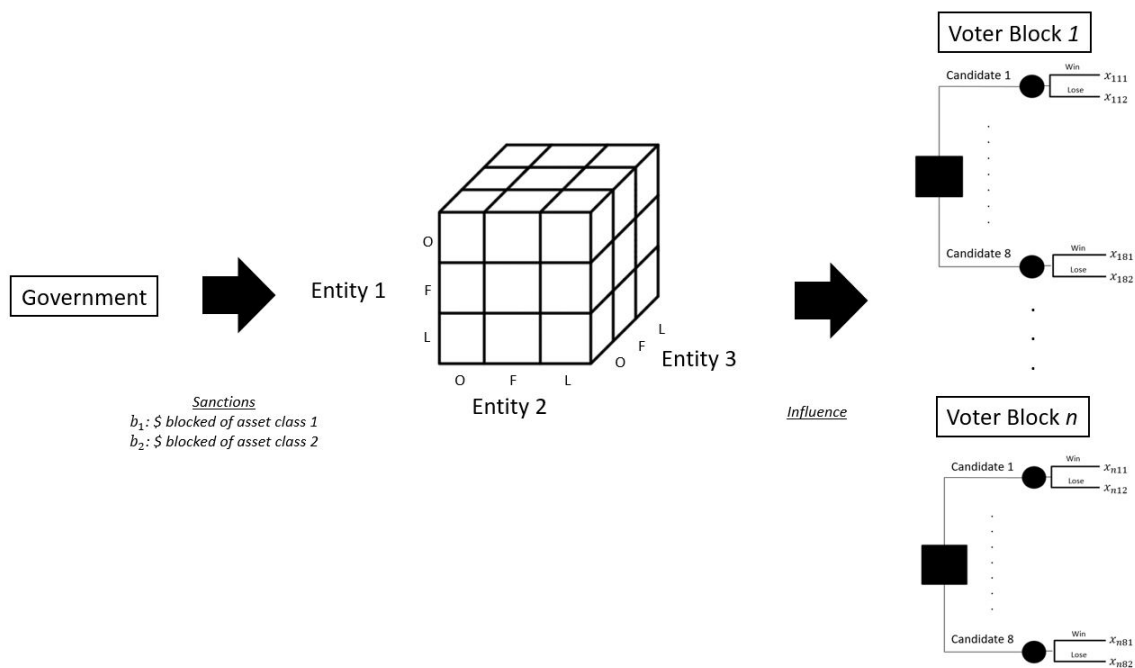


Figure 10. Detering Improper Electoral Interference as a Regulated Prospect Game

This example considers a political primary with eight candidates, six of whom conform to traditional party standards (i.e., Candidates 1 through 6) and two of whom represent a shift from the party’s conventional platform (i.e., Candidates 7 and 8). It is assumed that Candidate 7 is sympathetic to positions espoused by Entities 1 and 2, whereas Candidate 8 holds positions aligned with to Entity 3.

Voters are partitioned among several blocks that respectively aggregate and represent voters who behave in a similar manner (i.e., voters within a block have CPT

parameters, utilities and associated probabilities assumed to be identical). A decisionmaker is therefore a *representative* voter from the block that decides for whom to vote based upon the candidate’s appeal (x_{ijk}), which is interpreted as a gain or loss respective to the sitting president, and their subjective assessment of him or her winning (p_{ijk}). Whereas a voter could make a decision by explicitly considering the actions of all other voters, in a manner analogous to that discussed by Levy (2003), we model this scenario as a decision problem wherein a voter’s perception of the preferences of others informs their likelihood assessment that a candidate will win.

We model $n = 50$ blocks of voters with static CPT-parameter profiles (i.e., persuasion does not affect risk attitude). More specifically, we assume voter blocks evaluate the lotteries utilizing equations (19a)–(19d) in accordance with the randomly assigned parameters listed in Table 7, wherein the initial gain/loss values (\hat{x}_{ijk}) and their associated probabilities (\hat{p}_{ijk}) can also be found. Each voter’s CPT-parameters are sampled from a uniform distribution with a median estimate consistent with Tversky and Kahneman (1992), and the prospect-specific parameters (i.e, \hat{x}_{ijk} and \hat{p}_{ijk}) are established to allow heterogeneous initial voter perception. Since PGs and RPGs require these parameters to be known, we primarily focus on a single instantiation of the underlying PG in Sections 3.4 – 3.4 using the randomly generated point estimates. However, in Sections 3.4 and 3.4, we discuss how our models can inform further analysis should these parameters be uncertain.

Table 7. Uniform Distribution on voter CPT-parameters, and \hat{x}_{ijk} - and \hat{p}_{ijk} -values

α_i	β_i	λ_i	δ_i	σ_i	\hat{x}_{ijk}	\hat{p}_{ijk}
[0.80, 0.96]	[0.80, 0.96]	[2, 2.5]	[0.56, 0.66]	[0.64, 0.74]	[-1, 1]	[0, 1]

Each persuader may take one of three actions with respect to the election by engaging in an onslaught interference campaign (i.e., action O), illegal fringe operations (i.e., action F), or law-abiding activity (i.e., action L). The collective action of these

entities may affect each voters x_{ijk} - and p_{ijk} -values in the following manner:

$$x_{ijk} = \hat{x}_{ijk} + \sum_{p \in \mathcal{P}} t_{ijk}^{pl}, \quad \text{where } l = a_p$$

$$p_{ij1} = \hat{p}_{ij1} + \sum_{p \in \mathcal{P}} b_{ij1}^{pl}, \quad \text{where } l = a_p$$

$$p_{ij2} = 1 - p_{ij1}.$$

The parameters t_{ijk}^{pl} and b_{ij1}^{pl} are the effects that Entity p 's action l has on voter i 's perception of the value and probability, respectively, associated with the k -th outcome of voting for candidate j . Of note, the persuasion update for probabilities is defined explicitly with respect to $k = 1$ for notational clarity in observance of the axioms of probability. Each of the persuasion update parameters are randomly selected in accordance with a uniform distribution over the range $[-0.5, 0.5]$ for t_{ijk}^{pl} and $[(1 - \hat{p}_{ij1})/3, \hat{p}_{ij1}/3]$ for b_{ij1}^{pl} to allow for the possibility of meaningful persuasive action.

Through their persuasive actions, each Entity p tries to maximize their own objective. Entities 1 and 3 are each assumed as seeking to maximize the number of voters supporting their respective sympathetic candidate (i.e., Candidates 7 and 8 respectively), whereas Entity 2 desires to create tension in the electorate by maximizing support for both Candidates 7 and 8. Therefore, for any voter i , the corresponding q_{ip} for Entities 1 and 3 correspond with Candidates 7 and 8, respectively, but Entity 2 has q_{ip} corresponding to Candidate 7 for half of the voter blocks and Candidate 8 for the remaining blocks.

Persuader utility is also affected by punitive measures imposed by the government (i.e., the regulator). Mirroring the legislation proposed in the Deter Act in 2018 (Rubio and Van Hollen, 2018), this example considers possible regulating actions

against two classes of financial assets (i.e., b_1 and b_2) in millions of dollars. These sanctions affect each entity's payoff as represented in equations (26)–(28):

$$\text{Entity } p\text{'s Payoff (Onslaught): } -0.01b_1 - 0.02b_2 + \sum_{i \in \mathcal{I}} \phi_{ip}, \quad (26)$$

$$\text{Entity } p\text{'s Payoff (Fringe): } -0.01b_1 + \sum_{i \in \mathcal{I}} \phi_{ip}, \quad (27)$$

$$\text{Entity } p\text{'s Payoff (Legal): } \sum_{i \in \mathcal{I}} \phi_{ip}. \quad (28)$$

The monitoring capabilities of the regulator are assumed to be robust such that the true action taken by the persuaders can be detected. The coefficients assigned to a sanction's effect are notional but assumed to reflect the entity's value calculus in the mold of multi-objective decision analysis. For illustration purposes, this balance is assumed to be identical across all persuaders.

Table 8. Regulator Objective Function Coefficient Values

Entity	ξ_{pO}	ξ_{pF}	ξ_{pL}
1	-10	-2	1
2	-9	-3	1
3	-8	-1	1

The government would prefer all influence to occur in accordance with established legal standards. However, if some form of interference by a persuader cannot be prevented outright, the government's preference is for fringe activity as opposed to an aggressive, onslaught campaign. The negative effect associated with such interference is assumed to vary depending on the entity's capability and other political factors. These notional effects are listed as ξ_{pO} , ξ_{pF} and ξ_{pL} in Table 8. Utilizing these effects, the objective function of the government with respect to the Nash equilibrium or cognitive hierarchy settings, and the correlated equilibrium setting can be seen in equations (29) and (30), respectively. The two equations are conceptually identical

but adapted to the terminology of the respective modeling frameworks.

$$\max_{a \in \mathcal{A}} \sum_{a \in \mathcal{A}} \left(S_{1a_1} S_{2a_2} S_{3a_3} \right) \left(\sum_{p \in \mathcal{P}} \xi_{pa_p} \right) - 0.01b_1 - 0.02b_2 \quad (29)$$

$$\max_{a \in \mathcal{A}} \sum_{a \in \mathcal{A}} S'_a \left(\sum_{p \in \mathcal{P}} \xi_{pa_p} \right) - 0.01b_1 - 0.02b_2 \quad (30)$$

Finally, the government's decision variables b_1 and b_2 are both assumed to be bounded within $[\$0, \$60M]$, and instances are considered such that they may take on continuous values or discrete values on the order of millions of dollars.

Utilizing the models and methods described in Section 3.3, this RPG under the solution concepts of Nash equilibrium, correlated equilibrium and cognitive hierarchy are applied utilizing the global solver BARON on an HP Z820 equipped with a 2.60 GHz Intel E5-260 processor and 192GB of RAM.

Perfectly Rational Persuaders: Nash Equilibrium.

To understand whether intervention (i.e., regulation) is necessary, it is advisable to first analyze the PG generated with null regulation. In this example, a null action by the regulator would consist of no punitive sanctions being imposed for election interference.

		Entity 3		
		Onslaught Campaign	Fringe Operations	Legal Activity
		Entity 2		
		O	F	L
Entity 1	O	(5,5,6)	(5,5,9)	(5,5,9)
	F	(6,6,12)	(4,4,12)	(4,4,7)
	L	(3,3,8)	(3,3,8)	(4,4,7)
		Entity 2		
		O	F	L
Entity 1	O	(6,6,8)	(6,6,9)	(5,5,10)
	F	(7,7,12)	(4,4,9)	(4,4,8)
	L	(3,3,8)	(4,4,6)	(3,3,6)
		Entity 2		
		O	F	L
Entity 1	O	(5,5,8)	(5,5,9)	(5,5,11)
	F	(4,4,12)	(2,2,9)	(3,3,10)
	L	(3,3,8)	(4,4,6)	(3,3,5)

Figure 11. Induced Prospect Game with No Sanctions Imposed

The payoff structure of this PG can be seen in Figure 11 wherein the tuple in each cell represents the number of voting blocks supporting the preferred candidate

of Entity 1, Entity 2, and Entity 3, respectively. The optimistic-Nash equilibrium with respect to the government's objective function consists of the following pure strategies: Entity 1 conducts an onslaught campaign, and Entities 2 and 3 take lawful action. Under this persuader action profile, Candidate 2 garners the support of the most voting blocks at 12, whereas Candidate 8 is in a three-way tie for third and Candidate 7 is in sixth place. Furthermore, the government's objective function for this baseline has a value of -8. Under pessimistic assumptions such that the entities play the equilibrium that most negatively affects the regulator, Entities 1 and 3 conduct onslaught campaigns and Entity 2 engages in fringe operations such that the government's objective function equals -21, a lower bound on the optimal solution to the RPG.

Since this preliminary analysis finds interference will occur with certainty in the election, intervention strategies should be considered. Therefore, equations (26)–(29) are incorporated into the math program described in Section 3.3, and it is solved find an optimal \mathbf{b} .

The optimal solutions for when \mathcal{B} is a continuous space, or when it is a discrete space assuming integer-valued increments of millions of dollar, are very similar. In the discrete variant, the government selects $b_1=\$60M$ and $b_2=\$53M$, whereas in the continuous instance the government levies sanction of $b_1=\$60M$ and $b_2=\$53.109M$. These punitive actions result in the mixed strategy equilibriums seen in Table 9, wherein each entity's strategy vector lists the probability of conducting an onslaught campaign, fringe operations, or lawful activity, respectively. Given the relative proximity of these two solutions, only the discrete variant is discussed hereafter.

Under this solution, the regulator's objective function value with optimistic assumptions increases to -2.8. The results also imply there is a probability of 0.556 that all entities involved choose to engage in lawful activity, whereas without intervention

Table 9. Election Inference RPG - Optimal Optimistic-Nash Equilibriums

S_p	Discrete \mathcal{B}	Continuous \mathcal{B}
S_1	(0.098, 0.000, 0.902)	(0.097, 0.000, 0.903)
S_2	(0.000, 0.259, 0.741)	(0.000, 0.256, 0.744)
S_3	(0.168, 0.000, 0.832)	(0.168, 0.000, 0.832)

there is no possibility of such an action profile occurring. However, the probability that at least one entity will engage in an onslaught campaign is approximately 0.25. This result represents a reduction from the PG with null regulation wherein Entity 1 conducted an onslaught campaign with certainty. Therefore, the sanctions are a stabilizing influence, increasing the probability that the electoral process proceeds without unlawful interference. Furthermore, under pessimistic assumptions, the government's objective function value increases from -21 to -9.66. In fact, the entity's strategy profiles under the pessimistic-Nash equilibrium with $b_1=\$60M$ and $b_2=\$53M$ coincide with the optimistic-Nash equilibrium of the aforementioned PG: Entity 1 conducts an onslaught campaign, and Entity 2 and 3 take lawful action.

As noted previously, the government does not take any action to directly influence the voters. However, the sanctions to deter interference have second-order effects based upon how it changes the persuaders' behavior. These effects can be observed in Figure 12 wherein the probability distribution for each candidate of finishing the election in a given position is depicted.

Consider the government's ideal persuader action profile wherein all entities conduct operations in accordance with the law. Should this occur, the ultimate candidate ordering after voting yields the following: Candidate 2, Candidate 5, Candidate 4, Candidate 3 (tied for forth), Candidate 1 (tied for forth), Candidate 8, Candidate 6, and Candidate 7. This result differs slightly from the expected ordering of the candidates if the mean of each distribution from Figure 12 is utilized as a measure of central tendency. Namely, Candidate 2 still would be expected to win, but the

ordering of candidates in other positions changes (e.g., Candidate 5 is expected to finish third and Candidate 8 to finish fifth).




Figure 12. Distributions of Finishes for Optimistic-Optimal Nash Equilibrium RPG

These results also differ from the baseline PG wherein no sanctions are imposed. Under the optimistic-optimal solution of the RPG, Candidate 2 places first with a 0.60 probability instead of certainty; Candidate 8 places in the top two positions with a probability of 0.25 instead of a third-place tie with certainty; and Candidate 7 finishes in the bottom two positions with a probability of 0.90 instead of sixth with certainty. In maintaining the integrity of the election, the government actually *increased* the probability that one of the entity’s preferred opponents wins (i.e., Entity 3 favored Candidate 8). Superficially, such a result is counterintuitive; however, it is in fact consistent with the government’s goal of safeguarding the election’s integrity and not explicitly supporting any candidate.

Perfectly Rational Persuaders: Correlated Equilibrium.

As in the Nash equilibrium instance, the PG depicted in Figure 11 is evaluated first. In doing so, it is identified that the government's preferred correlated equilibrium is also the optimistic Nash equilibrium for the PG. Namely, with certainty Entity 1 conducts an onslaught campaign, and Entity 2 and 3 take lawful action.

Since Entity 1 will interfere with certainty in the election, intervention strategies should be considered. Thus, equations (26)–(29) augment the math program described in Section 3.3, and it is solved to find an optimal \mathbf{b} and S' .

The solutions for a discrete or continuous \mathcal{B} coincide exactly with the government instituting the maximum sanction values of $b_1 = b_2 = \$60M$. The associated correlated equilibrium is depicted in Table 10 wherein $\mathbf{a} \in \mathcal{A}$ is listed in order of the entities (i.e., event FOL coincides with Entity 1 fringe operations, Entity 2 onslaught campaign, and Entity 3 lawful activity). It can be observed that only 8 of the 27 possible collective actions are in the support, such that the remaining 19 have zero probability of being signaled.

Table 10. Election Inference RPG - Optimal Correlated Equilibriums

S'_a	Discrete \mathcal{B}	Continuous \mathcal{B}
S'_{LLL}	0.659	0.659
S'_{LLF}	0.044	0.044
S'_{LLO}	0.093	0.093
S'_{LFL}	0.082	0.082
S'_{FLL}	0.082	0.082
S'_{FFO}	0.010	0.010
S'_{FOF}	0.029	0.029

When compared to the RPG's optimistic Nash equilibrium solution, the government can expect a higher probability of combined lawful actions by all entities. To wit, in the Nash equilibrium variant, the government can expect such an event to occur with a probability of 0.556 versus 0.659 in the correlated equilibrium instance.

There also exists analogous reductions in the regulator's objective function value, and in the probability of one entity conducting an onslaught campaign. The optimal objective function value increases from -2.8 under the optimistic-optimal Nash equilibrium to -0.904 under the correlated equilibrium, and the probability of at least one entity conducting an onslaught campaign reduces from approximately 0.25 to 0.132. Such reductions are possible because all Nash equilibria are correlated equilibria (although the reverse does not hold true) and, as such, the optimal objective function of an RPG characterized by correlated equilibria is bounded below by the optimal objective function when it is characterized by Nash equilibria.

However, the practical difficulty of utilizing a correlated equilibrium in this context is that a regulator must signal publicly utilizing the distribution listed in Table 10, and a realization of this signal may induce interference. By doing so, the government is able to increase the probability of an interference-free election but at the expense of accepting that some persuader interference may occur. Politically, such an admission is problematic, yet should some regulating government find itself in a position of insurmountable vulnerability, such a solution may be acceptable, and the political effects can be potentially mitigated by signaling through ambiguous and/or conflicting statements.

As with the Nash equilibrium RPG, the solution generated by solving the correlated equilibrium RPG yields second-order effects on the election's outcome. These effects are similar to those found under the Nash equilibrium concept, but they generally exhibit less uncertainty. That is, the probability of a candidate's mode position from the Nash equilibrium RPG occurring increases across all candidate by reducing the probability of other outcomes (e.g., the probability of Candidate 2 finishing in first place increases from 0.60 to 0.785 by reducing the probability of a second and fourth place finish from 0.212 and 0.194 to 0.132 and 0.082, respectively). Likewise,

when comparing the expected election results under the correlated equilibrium to the results under the government’s preferred persuader action profile of collectively legal activity, there are only mild differences with regard to candidate ordering (i.e., the tie between Candidates 1 and 3 is broken in favor of Candidate 1). Such minor consequences are promising in that they indicate the government’s regulating action can be expected to produce minimal second-order effects on the electoral outcome.




Figure 13. Distributions of Finishes under Optimal Correlated Equilibrium RPG

Boundedly Rational Persuaders: Cognitive Hierarchy.

In modeling the bounded rationality of Entities 1–3, a value of $\hat{\tau}=1.5$ is utilized (Camerer et al., 2004), and this value is assumed to be static for demonstrative simplicity. As with the analysis under perfect rationality, we first consider the PG depicted in Figure 11. The resulting cognitive hierarchy solution profile can be seen in Table 11.

The cognitive hierarchy solution provides an expected distribution of selected strategies provided the mean-level of thinking depth coincides with τ . The resulting

Table 11. Election Inference PG - Cognitive Hierarchy Solution

S_p	Strategy Profiles
S_1	(0.599, 0.325, 0.074)
S_2	(0.850, 0.074, 0.074)
S_3	(0.409, 0.074, 0.516)

behavior for the sanction-free PG is unique when compared to the Nash equilibrium and correlated equilibrium variants. That is, in the cognitive hierarchy PG there exists a probability, albeit remote, of 0.0028 that all entities engage in legal activity, and there is a probability of 0.96 that at least one entity conducts an onslaught interference campaign. Conversely, the Nash equilibrium and correlated equilibrium PGs conclude that an onslaught campaign will occur with certainty. Setting these differences aside, the cognitive hierarchy PG still demonstrates a high probability of interference and, as such, the deterring effect of sanctions should be explored.

Table 12. Election Inference RPG - Cognitive Hierarchy Solutions

S_p	Discrete \mathcal{B}	Continuous \mathcal{B}
S_1	(0.111, 0.184, 0.705)	(0.111, 0.184, 0.705)
S_2	(0.075, 0.075, 0.850)	(0.075, 0.075, 0.850)
S_3	(0.075, 0.075, 0.850)	(0.075, 0.075, 0.850)

The resulting optimal solutions for both discrete and continuous \mathcal{B} are presented in Table 12 wherein $b_1 = b_2 = \$60M$. As with the correlated equilibrium RPG, these solutions coincide for the cognitive hierarchy RPG.

Under this solution, the regulator's objective function value increases to -2.43 and, similar to the other RPGs examined, the optimal solution of the cognitive hierarchy RPG necessitates sanctions near the upper end of their feasible regions. The optimal sanction values also coincide with the correlated equilibrium solution; however, the resulting behavior these sanctions induce is different. In the cognitive hierarchy RPG, these sanction values induce collectively law-abiding behavior in the persuaders with probability 0.509, and the probability that at least one entity engages in an on-

slaught campaign is 0.236. These results stand in contrast to the optimal correlated equilibrium RPG solution wherein the same sanction values induced probabilities of 0.659 and 0.132, respectively. Moreover, the cognitive hierarchy model produces the most pessimistic results regarding the possibility of all entities acting in accordance with the law, but project a higher value of the government's expected utility than the optimistic-Nash equilibrium RPG solution.




Figure 14. Distributions of Finishes under Optimal Cognitive Hierarchy RPG

With regard to the electoral outcome, the cognitive hierarchy RPG provides results that are more uncertain than the correlated equilibrium RPG and, depending on the specific candidate, are either more or less uncertain than the Nash equilibrium RPG. This outcome can be observed by noting the differences in the ordinal finish-related mode for all candidates between Figures 14, and either Figures 12 or 13. That is, the probability of the mode value occurring for all candidates is lower in the cognitive hierarchy RPG than the correlated equilibrium RPG; whereas, for Candidates 1 and 2, the cognitive hierarchy RPG solution yields a higher probability of the mode value than the Nash equilibrium RPG solution and a lower probability

for the remaining candidates. Moreover, when comparing the expected ordering of candidates in the cognitive hierarchy RPG, more differences are observed with respect to the government's preferred persuader action profile of collective legal activity than for the other two RPG solutions. The expected winner remains Candidate 2, but more drastic changes occur in the overall ordering of candidates (e.g., Candidate 8 has a mean finish of fourth place instead of sixth).

Uncertainty in the Prospect Game Payoff Structure.

Depending on the regulator's level of confidence in the point estimates utilized in the previous analyses (i.e., the decisionmakers' perceptions of the prospects and their CPT parameters), the regulator may wish to consider PGs and RPGs under differing assumptions. We illustrate the type of information that can be garnered via a suitable approach for this uncertainty by utilizing additional information from Table 7.

For a given set of input parameters, the models developed herein are deterministic. When considering parametric distributions, the regulator can gain insight into system behavior by sampling a sufficiently large set of the requisite parameters and analyzing each of the resulting PGs and RPGs. Space-filling designs are a particularly useful tool for such endeavors; however, for demonstrative clarity, we illustrate a regulator examining ten plausible scenarios generated randomly from the underlying distributions.

For each of these ten scenarios, the analyses conducted in Section 3.4–3.4 are repeated. Table 13 reports the results for each of four solution concepts: the Nash equilibrium (NE), alternatively under optimistic (O) and pessimistic (P) perspectives, the correlated equilibrium (CE), and cognitive hierarchy (CH). For each solution concept and scenario, tabulated are the objective function value and the respective

probabilities that at least one entity conducts an onslaught campaign or engages in fringe action.

Table 13. Solutions of PGs for 10 Additional Scenarios

Solution	Result	Scenario									
		1	2	3	4	5	6	7	8	9	10
NE	Objective (O)	-5.15	0	-1	-6	3	-7	-7	1	-10	0.33
	Prob Onslaught (O)	0.39	0	0	0	0	1	0.5	0	1	0
	Prob Fringe (O)	0.93	1	1	1	0	0	1	1	1	0.67
	Objective (P)	-23	-9	-7	-20	-21	-24	-21	-20	-17.33	-16
	Prob Onslaught (P)	1	1	1	1	1	1	1	1	1	1
	Prob Fringe (P)	0.5	1	0	1	0	0.5	1	1	1	0
CE	Objective	-3.7	0	-1	-2.79	3	-7	-7	1	-4.5	0.6
	Prob Onslaught	0.4	0	0	0.054	0	1	0.5	0	0.5	0
	Prob Fringe	0.8	1	1	0.51	0	0	1	1	0.5	0.67
CE	Objective	-9.2	-8.77	-1.67	-17.17	-5.21	-20.06	-16.02	-9.96	-12.31	-12.03
	Prob Onslaught	0.87	0.67	0.21	0.98	0.53	0.99	0.93	0.58	0.92	0.93
	Prob Fringe	0.87	0.81	0.58	0.87	0.42	0.58	0.61	0.99	0.58	0.42

Utilizing this information, the regulator can consider conditions under which intervention may be necessary. For instance, Scenario 1 yields a low objective function value across all solution concepts and results in the least objective function value under the pessimistic Nash solution. Analogous behavior can be observed for Scenarios 4, 7, and 9, implying that this subset of scenarios represent instances for which the regulator is most vulnerable.

Conversely, Scenario 3 yields a relatively high objective function value across all solution concepts and results in the greatest objective function value under the pessimistic Nash solution. Scenario 2 exhibits similar behavior and, collectively, they represent the least vulnerable instances for a regulator. However, Scenarios 5, 6, 8 and 10 exhibit no discernible pattern across game theoretic solutions, indicating that persuader behavior may be less predictable in these scenarios.

For a regulator examining each of the ten corresponding RPGs, Table 14 reports the optimal sanction values and the resulting persuader behavior for each of the scenarios and solution concepts. It can be observed that the NE and CE solutions coincide for the majority of the RPGs. Likewise, for the less vulnerable block of

RPGs (i.e., Scenarios 2 and 3) and three of the ambiguous RPGs (i.e., Scenarios 5, 8, and 10), the optimal RPG solution is to not intervene with sanctions. Under these scenarios, this implies that if the regulator believes the entities are perfectly rational, inaction is the regulator's optimal decision.

Table 14. Solutions of RPGs for 10 Additional Scenarios

Solution	Result	Scenario									
		1	2	3	4	5	6	7	8	9	10
NE	b_1	0	0	0	0	0	0	56.77	0	58.87	0
	b_2	50	0	0	50	0	50	21.62	0	0	0
	Objective (O)	-3	0	-1	0	3	0	-3	1	-8.05	0.33
	Prob Onslaught (O)	0	0	0	0	0	0	0	0	0.897	0
	Prob Fringe (O)	1	1	1	1	0	0.67	1	1	0.15	0.67
	Objective (P)	-19	-9	-7	-19	-21	-22	-7.08	-20	-12.59	-16
	Prob Onslaught (P)	1	1	1	1	1	1	0.691	1	1	1
Prob Fringe (P)	0	1	0	0	0	1	0	1	1	0	
CE	b_1	60	0	0	0	0	0	60	0	60	0
	b_2	26.67	0	0	50	0	50	35	0	42.86	0
	Objective	-1.83	0	-1	0	3	0	-2.97	1	-0.44	0.6
	Prob Onslaught	0	0	0	0	0	0	0.05	0	0.13	0
	Prob Fringe	0.71	1	1	1	0	0.67	0.92	1	0.14	0.67
CH	b_1	60	60	41.76	46.12	35.79	60	60	60	60	60
	b_2	42.97	0	0	0.6	9.88	60	48.74	0	60	51.61
	Objective	-2.4	-2.18	-0.37	-1.86	-0.46	-3.09	-3.83	-1.71	-2.66	-2.36
	Prob Onslaught	0.23	0.22	0.21	0.21	0.21	0.28	0.21	0.21	0.28	0.21
	Prob Fringe	0.43	0.69	0.21	0.64	0.21	0.40	0.68	0.65	0.26	0.55

Alternatively, the CH solutions suggest a more aggressive approach is required across all scenarios, even in the less vulnerable Scenarios 2 and 3. This result likely stems from the fact that the cognitive hierarchy solutions are more conservative than the optimistic Nash and correlated equilibrium solutions with regard to an entity engaging in an onslaught campaign. That is, across all scenarios except Scenario 9, the CH solution predicts a higher probability of one or more entities engaging in such a campaign. Given the consequences of onslaught campaigns (i.e., see Table 8), this probabilistic assessment yields lower objective function values for the CH solutions relative to the other solution concepts.

Finally, across all scenarios and solution concepts (except the optimistic NE and CE solutions for Scenario 5), the probabilities that at least one entity will attempt

to interfere in the election is non-zero, even with regulator intervention. Such insight indicates the regulator should focus efforts on deterring onslaught campaigns with the knowledge that some fringe activity is likely to persist.

Discussion.

The probability of interference and the expected candidate ordering results from Sections 3.4–3.4, combined with the behavior observed in Section 3.4, collectively indicate that deterring improper influence for this application with minimal second-order effects is more difficult when encountering boundedly rational rather than perfectly rational persuaders.

In Sections 3.4–3.4, each of the RPGs provides similar results with regard to the optimal sanction values. Each solution identifies that b_1 should be set at its upper bound of $\$60M$, and only the Nash equilibrium RPG solution designates a b_2 -value less than $\$60M$.

Should the Nash equilibrium PG be solved with $b_1=b_2=\$60M$, the government's optimistic objective function value decreases to -3.73, yielding a relative optimality gap of 33%. Its pessimistic objective function value decreases to -9.8 from -9.66 with optimal sanction values (i.e., a 1.4% reduction). Therefore, if the government is uncertain as to which game theoretic solution concept should be applied, by setting $b_1=b_2=\$60M$ they may guarantee a quality solution under any of the three solution constructs examined.

Conversely, if the regulator is not confident in the point estimates utilized to predict decisionmaker behavior in Sections 3.4–3.4, the results of Section 3.4 provide a method for continued analysis. PGs and RPGs do not provide definitive answers under uncertain information, but can inform decision making by exploring alternative scenarios. Likewise, their use in such settings can help identify candidate solutions

which can be further analyzed using other robust techniques.

3.5 Conclusions

In this research, we defined two new game theoretic frameworks, prospect games and regulated prospect games, designed to inform defensive national security action when multiple adversaries target a nation's citizenry and attempt to influence their decisions. Prospect games and regulated prospect games respectively model (a) the interactions of competing persuaders affecting a populace and (b) the actions of a regulating agent to alter such a framework. Such games may take many forms, depending on the nature of persuaders' interactions and assumptions about their rationality. Thus, their general forms were presented, in addition to the illustrated instantiation of three special cases concerning the Nash equilibrium, correlated equilibrium, and cognitive hierarchy solution concepts. Furthermore, the utility of these games to inform regulatory decisions was illustrated for a contemporary application concerning the deterrence of improper electoral interference.

In the models set forth herein, we require the decisionmakers' available prospects and methods of evaluation to be common knowledge. As such, they are most useful in situations wherein decisionmaker behavior is well-studied (e.g., elections or commerce). Promising areas for future research pertain to the relaxation of the common knowledge assumption and a more sophisticated treatment of uncertainty. Whereas this relaxation would allow for the solution of scenarios wherein the requisite information is incomplete, the problem reductions illustrated herein may no longer be applicable, and the games would need to be solved via some alternative game theoretic (e.g., Harsanyi, 1967; Aghassi and Bertsimas, 2006) or simulation-based (e.g., Lempert et al., 2006) technique. Moreover, other promising areas of future inquiry pertain to relaxing the assumption of a countable persuader action space and the

utilization of alternative, scenario-specific theories of descriptive choice.

Prospect games and regulated prospect games are, ultimately, designed to be decision support tools for regulating the activity of agents in an economy based upon persuasion. As such, there exists a wide array of potential application areas for which they can be utilized in commercial (e.g., Castañeda and Martinelli, 2018), international political (e.g., Pamp et al., 2018), and domestic regulatory (e.g., Fedeli et al., 2018; Brandt and Svendsen, 2018) settings. However, for the potential of these models to be realized, future research must attempt to quantify the effects of persuasive activity on a decisionmaker's CPT-parameters, and their subjective utility and probability estimates. In the interim, prospect games and regulated prospect games provide a rigorous and quantifiable framework upon which regulation can be constructed.

IV. Robust Influence Modeling under Structural and Parametric Uncertainty: An Afghan Counternarcotics Use Case

Abstract

An entity often wishes to influence the decisions of others in a system. This dynamic is apparent in a variety of settings including criminal justice, environmental regulation, and marketing applications. However, the central task of the influencing entity is confounded by uncertainty regarding their understanding of the structure and/or parameters of the decisions being made. The research herein sets forth a modeling framework to identify robust influence strategies under such uncertain conditions. Furthermore, the utility of this framework and its proper parameterization are illustrated via an application to the contemporary, global problem of the Afghan opium trade. Utilizing open source data, we demonstrate how counternarcotic policy can be informed using a quantitative analysis that embraces both the bounded rationality of the economy's decisionmakers and the government's uncertainty regarding the degree of this deviation from rationality. In this manner, we provide a new framework with which robust influence decisions can be identified under realistic information conditions, and we discuss how it can be used to inform real-world policy.

4.1 Introduction

Persuading an individual to adopt a given decision among a set of alternatives is an inherently difficult task. It requires a level of empathy to understand their full decision framework that may be difficult to achieve. In designing an influence strategy, a persuader must infer the answer to many questions: What other alternative prospects does the decisionmaker perceive? How does the decisionmaker evaluate this

set of prospects? How does the decisionmaker value each respective outcome? How does the decisionmaker understand the uncertainty giving rise to these outcomes? In what way will an influence action affect the decisionmaker and their perceptions?

The answers to such questions can be inferred via human-subject testing. However, the data collection requirements to do so with high confidence can be substantial, and the resulting statistical insights are likely only relevant to a specific context. If automated data collection efforts are suitable or if the underlying decision setting is predictable, such difficulties are less problematic; whereas the data collection requirements do not lessen, their assembly is facilitated. As the body of knowledge for applying statistical or machine learning techniques grows to answer the aforementioned questions, existent literature (i.e., Caballero et al., 2018) describes how an optimal influence strategy can be determined.

Unfortunately, there exist many situations for which automated data collection efforts are infeasible, or wherein the underlying decisionmaking setting is less predictable. If either of these conditions hold, then direct application of the models described by Caballero et al. (2018) is not possible; the parameter values required to formulate their *persuasion programs* are unknown. However, if bounded intervals for each these unknown parameters can be identified (i.e., an uncertainty set), alternative methods can be developed to generate robust influence actions that meet some threshold of performance regardless of the true parameter values. This research sets forth such a methodology by leveraging the models of Caballero et al. (2018) within the Robust Decision Making (RDM) framework set forth by Lempert et al. (2006), after which it demonstrates the methodology for a realistic use case.

The use case focuses on an application to a contemporary economic and national security problem: reducing supply for the Afghan opium poppy trade. This objective is part of an enduring counternarcotics effort in Afghanistan. The Afghan

illicit opium economy has long been suspected to be a primary revenue stream for the Taliban insurgency (SIGAR, 2018b), but development of a successful strategy against it has proven elusive (SIGAR, 2018a). To aid policy development, multiple economists have quantitatively modeled the effects of counternarcotic actions (e.g., Moreno-Sanchez et al., 2002; Clemens, 2008; Mejia and Restrepo, 2016); however, such models generally rely on equilibrium analysis that implicitly suggests the rationality of decisionmakers in the illicit economy. This research takes an alternative approach by considering the effect of counternarcotic strategies in a decision analytic instead of a game theoretic framework. By doing so, we demonstrate how a better understanding of a counternarcotic strategy's effect can be gained by embracing the bounded rationality of the illicit economy's decisionmakers and by acknowledging the associated uncertainty regarding their degree of departure from rationality.

The remainder of this chapter is structured as follows. Section 4.2 reviews the influence modeling paradigm, discusses the RDM framework, and introduces how the concepts can be leveraged jointly to inform influence actions under parametric and/or structural uncertainty. Section 4.3 illustrates the utility of our methodology via its application to the reduction of Afghan poppy cultivation in the Badakhshan province. Through this representative case study, we illustrate how policy decisions regarding influence can be quantitatively informed via our methodology. Finally, Section 4.4 discusses the implications of this research and avenues of future inquiry.

4.2 Robust Decision Making and Influence Modeling

The persuasion programs set forth by Caballero et al. (2018) pertain to the optimal manipulation of decision trees (e.g., see Figure 15). In a two-stage process, a persuader starts by influencing how a decisionmaker (or set of decisionmakers) perceives the underlying risk or uncertainty, values the payoff associated with each outcome, and/or

evaluates the available set of prospects. After this persuader action has been taken, the decisionmaker(s) selects a preferred prospect.




Figure 15. Example Decision Tree between Course of Actions (COAs) 1 and 2

Influence actions achieving such effects can take many forms. For instance, an appeal to the representativeness heuristic could alter subjective probabilities (Kahneman and Tversky, 1972), alternative frames could affect payoff valuations (Tversky and Kahneman, 1981), or emotional appeals could affect a decisionmaker’s risk attitude (Kugler et al., 2012). Other examples, in addition to a survey on the literature supporting such effects, are discussed in greater detail by Caballero et al. (2018).

Whereas the nature of the influence actions are assumed to be instance-specific, Caballero et al. (2018) describe their effects within the Cumulative Prospect Theory (CPT) framework (Tversky and Kahneman, 1992). CPT is an empirically-validated, descriptive theory of choice wherein an individual’s decision is described by utility valuations from a reference point; concave and convex utility for gains and losses, respectively; loss aversion; and systematic weighting of probabilistic information. More formally, consider a prospect f having n gain outcomes and m loss outcomes, respectively indexed as 1 through n and $-m$ through -1 . Each outcome value x_k measured from the reference point is evaluated *relatively* via CPT as

$$V(f) = V^+(f) + V^-(f), \text{ where} \quad (31a)$$

$$V^+(f) = \sum_{k=1}^n \pi_k^+ v(x_k), \quad (31b)$$

$$V^-(f) = \sum_{k=-m}^0 \pi_k^- v(x_k), \quad (31c)$$

$$\pi_n^+ = W^+(A_n), \quad (31d)$$

$$\pi_{-m}^- = W^+(A_{-m}), \quad (31e)$$

$$\pi_k^+ = W^+(A_k \cup \dots \cup A_n) - W^+(A_{k+1} \cup \dots \cup A_n), \quad 0 \leq k \leq n-1 \quad (31f)$$

$$\pi_k^- = W^+(A_{-m} \cup \dots \cup A_k) - W^+(A_{-m} \cup \dots \cup A_{k-1}), \quad 1-m \leq k \leq 0, \quad (31g)$$

wherein $v(\cdot)$ is a piecewise utility function concave for gains and convex for losses; W^+ and W^- are the event (outcome) weighting functions for gains and losses, respectively; and π^+ and π^- are the decision weights utilized to determine the respective component gains, $V^+(f)$, and component losses, $V^-(f)$.

The goal of a persuasion program is to maximize the number of decisionmakers selecting some preferred course of action by altering (1) the x_k -values, (2) the uncertainty associated with event A_k , and/or (3) the specific forms of the $W^\pm(\cdot)$ or $v(\cdot)$ functions. Accurately formulating a persuasion program requires a firm understanding of the structure of the decision tree, the baseline parameters (e.g., the decisionmaker's judged probabilities in the absence of influence), and how the persuader's action will alter the decisionmaker and their perceptions. Human subject testing (e.g. Booij et al., 2010; Campos-Vazquez and Cuijly, 2014; Schulreich et al., 2014) is one potential method to gain this understanding but, for a variety of reasons, it is not always practical and may not be feasible. Without these baseline parameters, a persuasion program cannot be utilized directly.

However, if uncertainty sets (or collections) can be identified for each unknown

factor, a robust influence strategy can be identified. Whereas the application of robust optimization is the intuitive approach to do so, herein we illustrate how the RDM framework set forth by Lempert et al. (2006) is preferable because of its flexibility.

Robust Decision Making.

RDM is an iterative decision analytic framework utilized in conditions of *deep uncertainty*. Deep uncertainty describes situations for which analysts cannot agree upon how to model the interaction of system variables, the underlying uncertainty, or the desirability of outcomes. RDM addresses these conditions by leveraging modern computer simulation, space-filling designs, and clustering algorithms.

To conduct an RDM analysis, it is necessary to identify the sets of available strategies, \vec{S} , and the future states of the world, \vec{F} . The objective of RDM is to select some $s \in \vec{S}$ that performs well across any future state of the world in \vec{F} . Lempert et al. (2006) propose the following five step iterative process to apply RDM: identify initial candidate robust strategies, identify vulnerabilities, suggest hedges against vulnerabilities, characterize deep uncertainties and trade-offs among strategies, and consider improved hedging options and surprises.

In practice, these steps consist of building a database $\vec{E} = \vec{S} \times \vec{F}$ (i.e., the futures ensemble) with some robustness metric assigned to each strategy-future combination. A candidate strategy s_{cand} is then selected, and a subset of future states in \vec{F} for which s_{cand} performs poorly is identified via clustering analysis. The next step is to determine alternative strategies less vulnerable to these future states but comparable to s_{cand} elsewhere, and select a new candidate strategy. This process is iteratively repeated until an acceptable course of action has been identified. After each round, there exists the potential to augment \vec{E} with more strategies or future states of the world as required.

Therefore, RDM requires \vec{F} to be a finite set. To account for this characteristic while accommodating situations wherein the set of plausible futures is excessively large (or infinite, as when defined on continuous domains of characteristic variables), a finitely-valued but adequately representative set \vec{F} can be generated via a space-filling design of the region. Numerous space-filling designs exist and provide various guarantees on how well the set of future states is explored.

Likewise, the effect obtained by utilizing some $s \in \vec{S}$ given a potential future state is determined by a scenario generator (i.e., a computer simulation). The use of a scenario generator, in combination with the set \vec{F} , allows for the modeling of various system variable interactions and their effect on system performance.

The raw values provided by the scenario generator are not the object of RDM analysis. Instead comparisons are made across strategies using some measure of robustness. Many alternative measures exist (e.g., absolute or relative regret), but the specific selection must ultimately be made in accordance with the problem setting (Roy, 2010). For example, if the degree of variation from the best-case is important, absolute measures may be appropriate. Alternatively, relative measures of robustness (or regret) can help highlight systematic differences. Likewise, the evaluation criteria utilized to distinguish across scenarios (e.g., minimize maximum regret, minimize expected regret, or minimize upper quartile regret) are also problem specific.

The conceptual steps in an RDM analysis are static. However, the space-filling design utilized to sample the true set of futures is contingent upon the underlying uncertainties. If all uncertainty sets are connected, then standard space-filling designs (e.g., a Latin hypercube) can be utilized; otherwise, less traditional designs capable of incorporating categorical variables (e.g., fast flexible designs) may be required. In either situation, multiple design candidates exist. Analysts must be cognizant that this design selection, in addition to the particular clustering analysis utilized to

identify vulnerabilities, may affect the final proposed strategy.

An RDM Approach to Influence.

In most influence settings, a persuader cannot be certain of many pertinent factors in the decision setting. The persuader can infer the prospects a decisionmaker(s) is considering but likely does not know these with certainty. Whereas the persuader confronts similar uncertainty with respect to a decisionmaker's judged probabilities, outcome valuations, risk attitude, loss aversion, *etc.*, it is likely that the persuader can bounded these factors within some range and subsequently apply the RDM framework to an influence setting.

As in Caballero et al. (2018), herein we assume a persuader is attempting to influence a set of decisionmakers, I , to each select some prospect under conditions of risk. The set of available influence strategies¹, \vec{S} , is known to the persuader. Each decisionmaker i faces a finite set of prospects J_i such that each $j \in J_i$ has a finite set of associated, uncertain outcomes, K_{ij} . That is, each decisionmaker is confronted with a decision tree akin to that depicted in Figure 15. The persuader is uncertain of the specific structural form of these decision trees, but we assume a persuader is able to infer a finite collection of decision trees such that each decisionmaker's true decision tree aligns with one of them. More formally, the persuader can identify the following uncertainty collections and uncertainty sets:

¹Caballero et al. (2018) describe this set as A , but herein we maintain consistency of notation with Lempert et al. (2006)

Decision Tree Uncertainty Collections

\mathcal{J}_i : A finite collection of prospect sets for decisionmaker i , one of which is the true J_i

\mathcal{K}_{ij} : A finite collection of outcome sets for decisionmaker i and prospect j , one of which is the true K_{ij}

Decision Tree Uncertainty Sets

\hat{Y}_{ijk} : Set of baseline raw values (\hat{y}_{ijk}) for the k^{th} event of prospect j for decisionmaker i

\hat{P}_{ijk} : Set of baseline probabilities (\hat{p}_{ijk}) for the k^{th} event of prospect j for decisionmaker i

These uncertainty collections and sets describe all potential decision trees the persuader believes the decisionmakers could consider. However, the manner in which the decisionmakers evaluate the decision trees is also a source of uncertainty, as is the effect of a persuader's influence action on both this evaluation and the decision trees themselves. Assuming the probability weighting functions and utility functions from Tversky and Kahneman (1992) are utilized (i.e., equations (32a)–(32c)), these sources of uncertainty can be described via the following uncertainty sets².

²We note that alternative probability weighting functions (e.g., Prelec, 1998) or utility functions (e.g., Wakker, 2010) can be substituted with relatively minor notational changes.

Decisionmaker Uncertainty Sets

$\hat{\Gamma}_i$: Set of baseline gain probability weighting coefficients ($\hat{\gamma}_i$)
for decisionmaker i

\hat{D}_i : Set of baseline loss probability weighting coefficients ($\hat{\delta}_i$)
for decisionmaker i

\hat{A}_i : Set of baseline gain utility coefficients ($\hat{\alpha}_i$) for decisionmaker i

\hat{B}_i : Set of baseline loss utility coefficients ($\hat{\beta}_i$) for decisionmaker i

$\hat{\Lambda}_i$: Set of baseline loss aversion ($\hat{\lambda}_i$) coefficients for decisionmaker i

\hat{R}_i : Set of baseline reference points (\hat{r}_i) for decisionmaker i

Influence Effect Uncertainty Sets

$F_{ijk}(s)$: Set of influence effect mappings ($f_{ijk}(s)$) on the baseline raw
value for the k^{th} event of prospect j for decisionmaker i

$G_{ijk}(s)$: Set of influence effect mappings ($g_{ijk}(s)$) on the baseline
probability for the k^{th} event of prospect j for decisionmaker i

H_i^θ : Set of influence effect mappings ($h_i^\theta(s)$) altering decisionmaker
 i 's baseline CPT-parameter $\hat{\theta}$ to θ , $\theta \in \{\gamma_i, \delta_i, \alpha_i, \beta_i, \lambda_i, r_i\}$

Moreover, this framework is also applicable to conditions of ambiguity per the results of Fox and Tversky (1998); however, select modifications are necessary to align the decision tree uncertainty sets with the tenets of Support Theory. That is, for each K_{ij} , an associated collection of sets \mathcal{C}_{ij} must be introduced, and the

uncertainty set \hat{P}_{ijk} must be replaced with $\hat{P}_{ij\Omega_u}$, respectively defined as follows:

\mathcal{C}_{ij} : Set of all non-empty subsets $\Omega_u \subseteq K_{ij}$,

$\hat{P}_{ij\Omega_u}$: Set of all baseline probabilities ($\hat{p}_{ij\Omega_u}$) of event disjunction Ω_u in prospect j by decisionmaker i .

Conceptually, the decision tree uncertainty collections (i.e., \mathcal{J}_i and \mathcal{K}_{ij}) describe the structural uncertainty, and the remaining uncertainty sets describe the parametric uncertainty. Collectively, they also constitute all known potential future states of the system. To ensure flexibility, we have left their form general; however, when discerning how to describe \vec{F} , the specific properties of each uncertainty is paramount. If all of the aforementioned collections are countable, then \vec{F} could potentially be taken via their enumeration. Otherwise, varying forms of space-filling designs should be considered to form \vec{F} . If \mathcal{J}_i or \mathcal{K}_{ij} has a cardinality greater than 1, then standard space-filling designs such as the Latin hypercube are not applicable. From the perspective of experimental design, each possible J_i or K_{ij} is a categorical factor level. Therefore, alternative designs such as a sliced Latin hypercube (Qian, 2012; Ba et al., 2015) should be considered. Other complicating factors restricting which space-filling designs can be utilized relate to the probabilistic nature of the uncertainty. When influencing a decisionmaker under risk, the axioms of probabilities must be satisfied, implying that the underlying design space is constrained. In such instances, the fast flexible design by Lekivetz and Jones (2015) is an attractive alternative.

Once \vec{F} has been determined, the value of each element in the futures ensemble \vec{E} can be identified utilizing the scenario generator. In modeling influence, the scenario generator consists of a relatively direct application of CPT. Because each element of \vec{E} considers a particular future, a specific element of each aforementioned uncertainty

collections or sets is provided. In turn, each decision tree and its respective decisionmaker's evaluation calculus is fully specified in its post influence state. As such, once the outcomes for each prospect have been sorted and partitioned into gains and losses, equations (31a)–(31g) can be applied to determine which prospect decisionmaker i prefers. More formally, for each decisionmaker i and a specified set of available prospects and associated outcomes (i.e., $J_i \in \mathcal{J}_i$ and $K_{ij} \in \mathcal{K}_{ij}$), a set of outcome values and judged probabilities (i.e., y_{ijk} and p_{ijk}), and a tuple of CPT-parameters $(\gamma_i, \delta_i, \alpha_i, \beta_i, \lambda_i, r_i)$ is defined by

$$\begin{aligned} y_{ijk} &= \hat{y}_{ijk} + f_{ijk}(s), \quad \forall i \in I, j \in J_i, k \in K_{ij}, \\ x_{ijk} &= y_{ijk} - r_i, \quad \forall i \in I, j \in J_i, k \in K_{ij}, \\ p_{ijk} &= \hat{p}_{ijk} + g_{ijk}(s), \quad \forall i \in I, j \in J_i, k \in K_{ij}, \\ \theta_i &= \hat{\theta}_i + h_i^\theta(s), \quad \forall i \in I, \theta = \{\gamma_i, \delta_i, \alpha_i, \beta_i, \lambda_i, r_i\}, \end{aligned}$$

and

$$W^+(p_{ijk}) = \frac{(p_{ijk})^{\gamma_i}}{((p_{ijk})^{\gamma_i} + (1 - p_{ijk})^{\gamma_i})^{\gamma_i^{-1}}}, \quad (32a)$$

$$W^-(p_{ijk}) = \frac{(p_{ijk})^{\delta_i}}{((p_{ijk})^{\delta_i} + (1 - p_{ijk})^{\delta_i})^{\delta_i^{-1}}}, \quad (32b)$$

$$v(x_{ijk}) = \begin{cases} (x_{ijk})^{\alpha_i}, & x_{ijk} \geq 0 \\ -\lambda_i(x_{ijk})^{\beta_i}, & x_{ijk} < 0. \end{cases} \quad (32c)$$

These values are utilized in equations (31a)–(31g) to determine which prospect has the maximum value and is preferred by the decisionmaker.

The persuader's utility is in turn affected by the collective choices of the decisionmakers. It may be the case that, for each decisionmaker and prospect set J_i , the persuader prefers the decisionmaker to select some prospect $q(J_i)$, in a binary manner analogous to Caballero et al. (2018), or the persuader may value the decisionmaker's selection along some continuum. Therefore, the selection of this *scenario generator measure* is ultimately instance specific and depends upon the persuader's goal for the system. The persuader also must determine whether the successful influence of a plurality of decisionmakers is preferred or, conversely, if each decisionmaker's choice is weighted differently.

This selection of a scenario generator measure, in addition to the chosen robustness measure, evaluation criteria, and clustering algorithm are RDM tailorable components. Each must be selected in accordance with the persuader's needs and, realistically, software availability. As with RDM in a generic setting, the selection of the robustness measure and an evaluation criteria must be informed by the persuader's objective. For example, if deviations in natural units from optimal solutions in future states are important, then the use of expected absolute regret is a well-suited robustness measure. However, such an objective does allow for outliers to influence a decision and a persuader may alternatively seek to minimize the expected relative regret if this property is undesirable.

4.3 Case Study: Influencing Landowning Household Crop Choice in Badakhshan Province

Since the 1980s, opium poppy cultivation has steadily increased in Afghanistan (Ward and Byrd, 2004). Production soared when the Taliban took control in 1996 (Woody, 2018a), and today it has become an integral part of the rural Afghan economy (Byrd, 2017). In an unsecure and volatile environment, opium poppy cultivation

is an appealing choice for Afghan farmers due to its high selling price, long shelf life, and profitable byproducts that can be used for, e.g., heating and livestock feed (Mansfield and Fishstein, 2016). The extent of poppy production is so vast that Afghanistan has become the world's largest producer (Woody, 2017, 2018a). Likewise, the opium sector as a whole, including the licit and illicit activities it supports, is a major source of revenue for the country; some experts estimate drugs constitute between a third and a half of the overall Afghan economy (Felbab-Brown, 2016).

While the extent of opium poppy's effect on counterinsurgency is contested (Mansfield, 2018), much of the revenue associated with its trade is suspected to support the Taliban (SIGAR, 2018b). For this reason, counternarcotics and counterinsurgency have often been viewed as complementary campaigns by the United States. Unfortunately, no policy enacted over the course of the conflict has been able to meaningfully cripple the Afghan opium economy (SIGAR, 2018a). These failures are often attributed to a policy's poor underpinnings to economic, social and cultural realities (Mansfield and Fishstein, 2016). Whereas strategies considering such factors have long been advocated by Afghan experts (e.g., Ward and Byrd, 2004; Ward et al., 2008), the implementation of an effective counternarcotics policy has proven elusive (SIGAR, 2018a).

In this section, we illustrate how our modeling paradigm can be used to formulate policy informed by socio-economic realities and aimed at achieving a portion of the larger counternarcotics objective. Utilizing a variety of open source data, we formulate and solve a representative use case aimed at deterring the farmers in the northeastern province Badakhshan (Figure 16) from engaging in poppy cultivation. We model this province because it has witnessed high levels of poppy cultivation in areas under government control (UNODC, 2018a) and, as discussed by SIGAR (2018a), a counternarcotic strategy requires such control to be effective.



Figure 16. Badakhshan Province, Afghanistan

Seasonal Crop Decision by Badakhshan Landowners.

In this section, we illustrate how the structure of a farmer's decision regarding what crop to plant and cultivate in a given season can be inferred. Although we utilize a variety of open source data over a variable time frame to inform this use case, an official policy development would be best served with recent data, tailored to the specific area of interest. As such, what follows is a representation of influence modeling under structural and parametric uncertainty. It is not meant to promote a specific policy, but instead to demonstrate the utility of the proposed methodology.

Landowners in Afghanistan generally have small holdings (UNODC, 2003). The same observation holds in the province of Badakhshan (Pain, 2010). The Central Statistics Organization of Afghanistan (2013) categorized Badakhshan farmers into three demographics: irrigated, rain-fed, and garden-plot farmers. On average, individual members of these groups own 4.2 jeribs (0.84 hectare), 6.4 jeribs (1.28 hectares), and 1.1 jeribs (0.22 hectares) of land, respectively (Central Statistics Organization of Afghanistan, 2013). Likewise, each demographic respectively constitutes 36%, 45%, and 19% of the 135,000 landowning households in the province (Central Statistics

Organization of Afghanistan, 2013). The irrigated farmer demographic can be further subdivided based upon their irrigation system. A variety of irrigation systems are used in Afghanistan, including shallow wells known as *arhads* and man-made underground channels from aquifers known as *karizs*. Both of these systems are drought-resistant; however, other more vulnerable systems based on rivers, springs, and snow-melt also exist (Qureshi, 2002). With this understanding, the irrigated farmer demographic can be subdivided into drought-resistant irrigated (DRI), and drought-vulnerable irrigated (DVI) farmers. According to the Central Statistics Organization of Afghanistan (2007), between 35% and 45% of Badakhshan irrigated farmers have drought-resistant systems, compared to the Afghanistan-wide estimate of 24% (Central Statistics Organization of Afghanistan, 2018). Therefore, in this case study we model a “representative” decisionmaker from each of the four demographics. That is,

$$I = \{\text{DRI, DVI, RF, GP}\}.$$

The relative distribution of each demographic is assumed to be 13%, 23%, 45%, and 19% of the total landowning population in Badakhshan in accordance with lower bound drought-resistant estimate provided by Central Statistics Organization of Afghanistan (2007). Likewise, we henceforth assume both rain-fed and garden-plot farmers are vulnerable to drought conditions, and that a representative farmer from each demographic owns the aforementioned average landholdings.

The crops available to Badakhshan farmers vary by district. However, there are a few staple options (i.e., potato, wheat, and opium poppies) that are grown in every district, excluding Shiki (Central Statistics Organization of Afghanistan, 2007). Other options such as onions, tomatoes, and maize are also widely cultivated. Ideally, the

influence model would consider all crop combination considerations; however, due to open source data availability and the purpose of illustrating the RDM modeling and solution methodologies, we limit the crops available to farmers as a combination between wheat, potatoes, tomatoes, and opium poppies. To build J_i , we note that according to the Central Statistics Organization of Afghanistan (2018), the majority of Afghan households cultivate one or two crops, and a minority plant up to three. We assume each farmer demographic i has the option to cultivate their land as one of the elements in the set J_i defined in equation (33).

$$J_i = \left[\begin{array}{l} \text{Plant Opium Poppies (100\% of land)} \\ \text{Plant Wheat (100\% of land)} \\ \text{Plant Crop X (100\% of land)} \\ \text{Plant Opium Poppies, Wheat (50\%, 50\% of land each)} \\ \text{Plant Opium Poppies, Wheat (25\%, 75\% of land)} \\ \text{Plant Opium Poppies, Wheat (75\%, 25\% of land)} \\ \vdots \\ \text{Plant Opium, Wheat, Crop X (33\% of land each)} \\ \text{Plant Opium, Wheat, Crop X (50\%, 25\%, 25\% of land)} \\ \text{Plant Opium, Wheat, Crop X (25\%, 50\%, 25\% of land)} \\ \text{Plant Opium, Wheat, Crop X (25\%, 25\%, 50\% of land).} \end{array} \right] \quad (33)$$

By the definition of J_i , we assume a representative basket of available options for each demographic includes three crops. Due to their widespread cultivation, we assume opium poppies and wheat are part of this basket. However, we allow for uncertainty in the third crop option (i.e., Crop X) which can be either potatoes or tomatoes. The uncertainty collection \mathcal{J}_i therefore has a cardinality of two for each farmer demographic

i , wherein the definition in equation (33) is repeated for both possibilities of Crop X. Moreover, due to the small acreage of landholdings, we assume that the allocation strategies of 50-50 percent or 25-75 percent when planting two crops, and 33-33-33 percent or 50-25-25 percent when planting three crops are sufficient to approximate actual behavior, providing 16 available prospects. Further behavioral experimentation could be considered to ascertain how the CPT-editing phase affects J_i to reduce its size, but we refrain from doing so for this case study.

For this illustration, we assume a farmer considers two underlying uncertainties that affect their crop yield, and that for any decisionmaker i and prospect j , \mathcal{K}_{ij} is a singleton collection composed of the uncertainties defined in either equation (34) or (35). The amount of rainfall in a given season is a primary concern, and, if planting poppy, so is the possibility of an eradication raid by the government. Therefore, the uncertainty associated with a planting decision including poppy is

$$K_{ij} = \begin{bmatrix} \text{Ample Rainfall, No Raid} \\ \text{Drought, No Raid} \\ \text{Ample Rainfall, Raid} \\ \text{Drought, Raid} \end{bmatrix}, \quad (34)$$

and the uncertainty associated with a planting decision excluding poppy is

$$K_{ij} = \begin{bmatrix} \text{Ample Rainfall} \\ \text{Drought} \end{bmatrix}. \quad (35)$$

Having defined the decisionmakers and the decision tree uncertainty collections, the general form of each farmer's crop selection decision is depicted in Figure 17.

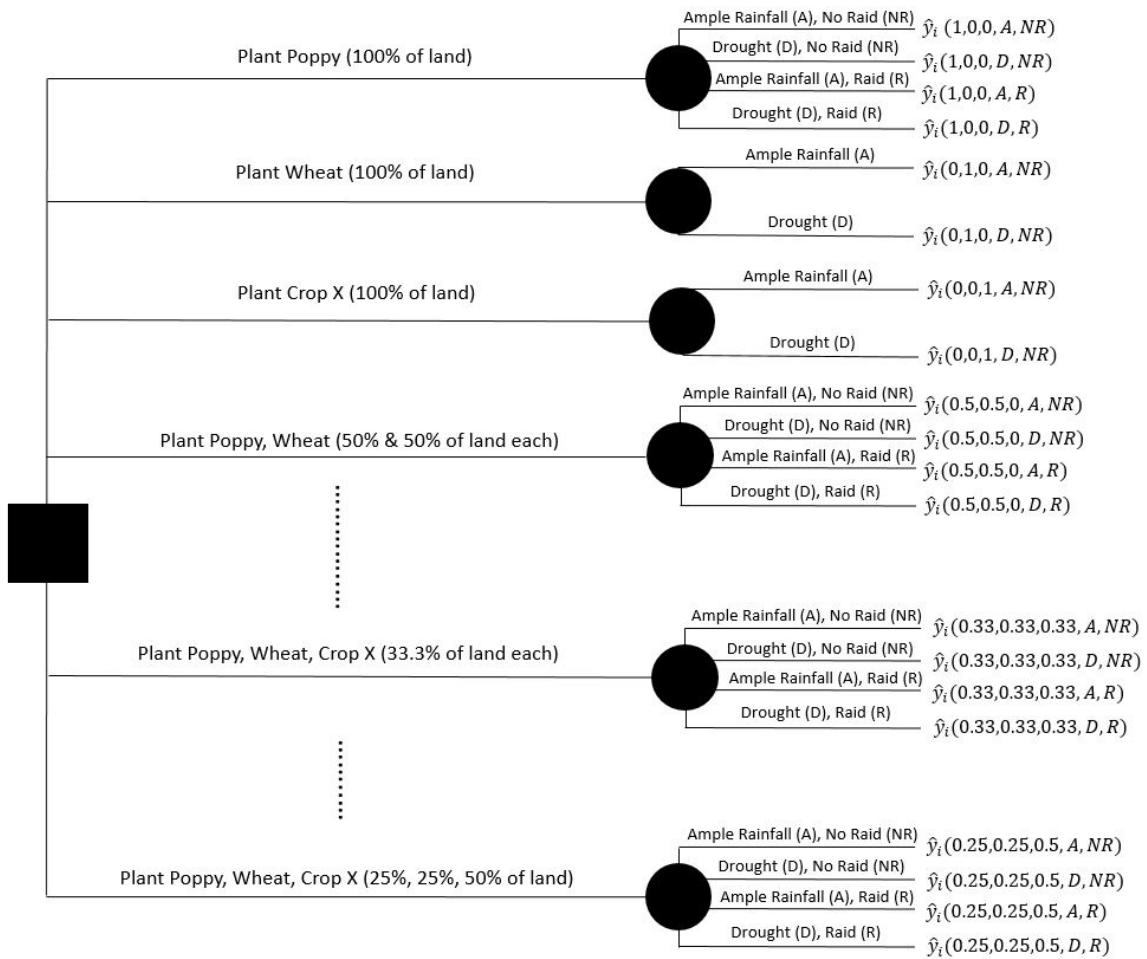


Figure 17. Generic Badakhshan Farmer Decision Tree

We introduce the following alternative notation to improve readability of this use case. The payoff (i.e., net income) associated with each prospect-outcome combination before influence (i.e., \hat{y}_{ijk}) can also be denoted as \hat{y}_i (% poppy, % wheat, % crop X, Rainfall Outcome, Raid Outcome). In the associated tuple, the first three elements describe the selected prospect j , whereas the final two elements describe the associated uncertain outcome k . A similar notational simplification is implemented for \hat{p}_{ijk} which can be written as \hat{p}_i (% poppy, % wheat, % crop X, Rainfall Outcome, Raid Outcome).

The uncertainty set \hat{Y}_{ijk} is compiled by combining the gross income (e.g., production per hectare, farm-gate price, etc.) and cost (e.g., *ushr*, labor, seed) estimates by Kuhn (2010) and from drought effects estimated by the UNODC (2018b). Kuhn (2010) provides net income projections per hectare for a variety of Afghan crops; however, these estimates were generated under the assumption of ample rainfall, judging by the coincidence of the poppy yield estimate with the actual yields observed in the 2017 bumper crop season (UNODC, 2018b). As such, the estimates listed by Kuhn (2010) are used as a baseline for a season with ample rainfall, and they are adjusted based on observed data from the UNODC (2018b) to obtain net income estimates in drought seasons. That is, we utilize the UNODC (2018b) data to inform poppy income estimates and extrapolate the observed trends to other crops. Although this technique is not ideal for modeling accuracy, we utilize this naive approach due to a lack of data availability. As such, the resulting uncertainty sets \hat{Y}_{ijk} for each demographic are formed by multiplying the respective lower and upper bounds of jerib net income by the number of jeribs utilized for each crop based on the given prospect j and outcome k . The per jerib income values are listed in Table 15. Note that labor costs are assumed to be static at the rate listed by Kuhn (2010), who does not itemize the price of irrigation maintenance.

Table 15. Estimated Per Jerib Net Income Range

Label	Farmer Demographic	Crop	Uncertain Outcome	Net Income per Jerib
u_1	DRI	Poppy	A, NR	[\$442, \$1492]
u_2	DRI	Poppy	D, NR	[\$342, u_1]
u_3	DRI	Poppy	A/D, R	-\$270
u_4	DRI	Wheat	A, NR	[-\$1, \$149]
u_5	DRI	Wheat	D, NR	[-\$15, u_4]
u_6	DRI	Potato	A, NR	[\$130, \$919]
u_7	DRI	Potato	D, NR	[\$55, u_6]
u_8	DRI	Tomato	A, NR	[\$74, \$377]
u_9	DRI	Tomato	D, NR	[\$45, u_8]
u_{10}	DVI	Poppy	A, NR	[\$442, \$1492]
u_{11}	DVI	Poppy	D, NR	[\$202, u_{10}]
u_{12}	DVI	Poppy	A/D, R	-\$270
u_{13}	DVI	Wheat	A, NR	[-\$1, \$149]
u_{14}	DVI	Wheat	D, NR	[-\$35, u_{13}]
u_{15}	DVI	Potato	A, NR	[\$130, \$910]
u_{16}	DVI	Potato	D, NR	[-\$50, u_{15}]
u_{17}	DVI	Tomato	A, NR	[\$74, \$377]
u_{18}	DVI	Tomato	D, NR	[\$5, u_{17}]
u_{19}	RF	Poppy	A, NR	[\$380, \$1430]
u_{20}	RF	Poppy	D, NR	[\$140, u_{19}]
u_{21}	RF	Poppy	A/D, R	-\$333
u_{22}	RF	Wheat	A, NR	[-\$1, 149]
u_{23}	RF	Wheat	D, NR	[-\$35, u_{22}]
u_{24}	RF	Potato	A, NR	[\$130, \$910]
u_{25}	RF	Potato	D, NR	[-\$50, u_{24}]
u_{26}	RF	Tomato	A, NR	[\$74, \$377]
u_{27}	RF	Tomato	D, NR	[\$5, u_{26}]
u_{28}	GP	Poppy	A, NR	[\$504, \$1554]
u_{29}	GP	Poppy	D, NR	[\$264, u_{28}]
u_{30}	GP	Poppy	A/D, R	-\$208
u_{31}	GP	Wheat	A, NR	[\$26, \$177]
u_{32}	GP	Wheat	D, NR	[-\$8, u_{31}]
u_{33}	GP	Potato	A, NR	[\$174, \$964]
u_{34}	GP	Potato	D, NR	[-\$6, u_{33}]
u_{35}	GP	Tomato	A, NR	[\$117, \$421]
u_{36}	GP	Tomato	D, NR	[\$48, u_{35}]

The net income estimates per jerib are disparate across demographics. For instance, considering that the average Badakhshan household consists of nine individuals (Nicolle, 2010), and that poppy harvesting is a culturally acceptable activity in which Afghan women partake (UNODC, 2003), the garden-plot demographic can accommodate labor requirements internally, even during the approximately two week, labor-demanding poppy harvest. Conversely, rain-fed farmers with larger land holdings are, at times, forced to hire contract labor. Analogous differences can also be observed in the disparate effects drought is assumed to have on each demographic.

For illustration purposes, we assume farmers make their decisions under conditions of risk. Likewise, since we assume independence between rainfall and poppy eradication raids, the probability associated with any outcome is simply the product of the associated rainfall and raid probabilities. A conservative approach is taken regarding the probability of drought, and it is assumed to be between [0.25, 0.75]. In accordance with the infrequent poppy eradication raids in Badakhshan in recent years (UNODC, 2018b), we assume they are viewed as improbable events with a range of [0.02, 0.05]. As such, for a prospect j with respective acreage allocations a_W and a_X for wheat and crop X without poppy cultivation, we have

$$\hat{P}_{ijk} = \left\{ (p_1, p_2) \left| \begin{array}{l} 0.25 \leq p_1 \leq 0.75, \\ 0.25 \leq p_2 \leq 0.75, \\ p_1 + p_2 = 1 \end{array} \right. \right\}$$

where

$$p_1 = p(\text{ample rain}) = p_i(0, a_W, a_X, A, \text{NR}), \text{ and}$$

$$p_2 = p(\text{drought}) = p_i(0, a_W, a_X, D, \text{NR}).$$

For a prospect j including poppy cultivation with an acreage percentage a_P , we have

$$\hat{P}_{ijk} = \left\{ (p_1, p_2, p_3, p_4) \left| \begin{array}{l} 0.25 \leq p_1 \leq 0.75, 0.25 \leq p_2 \leq 0.75, \\ 0.02 \leq p_3 \leq 0.05, 0.02 \leq p_4 \leq 0.05 \\ p_1 + p_2 = 1, p_3 + p_4 = 1 \end{array} \right. \right\}$$

where

$$p_1 = p(\text{ample rain}) = p_1,$$

$$p_2 = p(\text{drought}) = 1 - p_1,$$

$$p_3 = p(\text{raid}),$$

$$p_4 = p(\text{no raid}) = 1 - p_3,$$

and

$$p_i(a_P, a_W, a_X, A, \text{NR}) = p_1 p_4,$$

$$\vdots$$

$$p_i(a_P, a_W, a_X, D, R) = p_2 p_3.$$

Finally, we turn to our attention to the parameterization of the decisionmaker uncertainty sets. We leverage the works of Booij and Van de Kuilen (2009), Abdellaoui (2000), and Abdellaoui et al. (2007) to inform the uncertainty sets associated with the probability weighting, utility curvature, and loss aversion coefficients, respectively. In the case of the utility curvature coefficients, the uncertainty sets are taken to be the range of estimates from the literature provided by Booij and Van de Kuilen (2009), whereas the experimental results from Abdellaoui (2000); Abdellaoui et al. (2007)

directly inform the probability weighting and loss aversion uncertainty sets. A more simplistic approach is taken with the reference point uncertainty set, and it is derived from Table 15 by the range formed via the product of the minimum and maximum unit net income values with each demographic's assumed landholdings. The resulting sets are quantified in Table 16 wherein L_i refers to the assumed size of demographic i 's landholdings. Of note, demographic decisionmaker uncertainty sets are assumed to be identical (except for the scale of \hat{R}_i) due to the high degree of ambiguity relating to the manner in which Badakhshan farmers distinguish between prospects. To our knowledge, no CPT-related studies have been conducted on this population to inform more accurate model parameters.

Table 16. Decisionmaker Uncertainty Sets

$\hat{\Gamma}_i$	\hat{D}_i	\hat{A}_i	\hat{B}_i	$\hat{\Lambda}_i$	\hat{R}_I
[0.492, 0.708]	[0.588, 0.812]	[0.22, 1.01]	[0.61, 1.06]	[0.74, 8.27]	$[-333L_i, 1554L_i]$

Influence Actions and their Effects.

In the Afghan conflict, multiple counternarcotic strategies have been attempted. The influence actions these strategies rely upon can be broadly classified into four categories (SIGAR, 2018a): eradication, interdiction, alternative development, and political support mobilization. Eradication efforts focus on the destruction of a standing poppy crop, whereas interdiction efforts, to include narcotics seizures and the destruction of narcotics production facilities, are applied further along the supply chain. Alternative development programs can be viewed as positive reinforcement relative to the negative reinforcement of eradication. They are designed to reduce poppy cultivation by increasing the attractiveness of licit livelihood opportunities. Finally, political support mobilization is a wide ranging category including public awareness campaigns and other efforts designed to strengthen Afghan institutions. Herein, we

assume that the persuader's task is to select some combination of these four actions to reduce poppy cultivation in Badakhshan.

The effects of these influence actions is highly dependent of the specific form of their implementation. For example, the alternative development program in the Helmand Food Zone (HFZ) which focused on substituting wheat for opium poppy has been criticized as ineffective and ill-advised (Mansfield and Fishstein, 2016). However, alternative programs to build an enduring capacity with more profitable crops have been viewed favorably (SIGAR, 2018a). A similar debate relating to the form of eradication efforts (e.g., manual vs. aerial spraying) dominated much of the early stages of the Afghan conflict (Coll, 2019).

For this case study, we adopt a strategic-level perspective and describe how each influence effort can be expected to affect the decisionmaker and the decision setting. Eradication campaigns, regardless of their form, target the farmers themselves. As such, an increase in their intensity will necessarily coincide with an increase in the probability of a poppy farming raid, and the reduced overall supply will result in increased expected income for successful harvests (Martin and Symansky, 2006). Conversely, interdiction efforts have an indirect effect on farmers. By targeting upstream entities, they serve to reduce the demand observed by the poppy farmers (Martin and Symansky, 2006) and, in isolation, can be expected to decrease farm-gate values.

The effects of an alternative development program are varied. However, herein we assume that such programs targeted at potential poppy farmers are designed to increase the attractiveness of cultivating wheat, potato, or tomato via either reducing costs or increasing profits. Likewise, the effects of political support mobilization are diverse, but herein we assume such action is composed of a public awareness campaign and provincial leader engagement with the following effects. The public awareness campaign reinforces the criminality of poppy cultivation and underscores

the financial risks associated with a destroyed crop, whereas the provincial leaders are encouraged to increase poppy related taxes. As such, political support mobilization can be expected to increase a farmer's judged probability of a raid and decrease the net income associated with poppy.

Table 17. Actions Included (Y) and Excluded (N) in each Influence Strategy

Influence Strategy (s_{cand})	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}	s_{11}	s_{12}	s_{13}	s_{14}	s_{15}	s_{16}
Eradication	N	Y	N	N	N	Y	Y	Y	N	N	N	Y	Y	Y	N	Y
Interdiction	N	N	Y	N	N	Y	N	N	Y	Y	N	Y	Y	N	Y	Y
Alternative Development	N	N	N	Y	N	N	Y	N	Y	N	Y	Y	N	Y	Y	Y
Pol. Sup. Mobilization	N	N	N	N	Y	N	N	Y	N	Y	Y	N	Y	Y	Y	Y

All available influence strategies (i.e., \vec{S}) are listed in Table 17. Likewise, the uncertainty sets for influence effects that account for these dynamics are displayed in Table 18. For this case study, these uncertainty sets are assumed to be static across all decisionmakers. Each strategy, if employed, is assumed to have an additive effect on the associated parameter as listed. For example, for some decisionmaker i and prospect j (including poppy) wherein outcome k represents the ample rainfall, raid outcome, we have

$$G_{ijk}(s_2) = [0, 0.7],$$

$$p_{ijk} = \hat{p}_{ijk} + g_{ijk}(s_2), \quad g_{ijk}(s_2) \in G_{ijk}(s_2).$$

Moreover, if an influence strategy employs multiple actions simultaneously, the cumulative effect is assumed to be additive. For instance, if influence strategy s_8 is utilized instead of s_2 we would have

$$g_{ijk}(s_8) = g_{ijk}(s_2) + g_{ijk}(s_5), \quad g_{ijk}(s_1) \in G_{ijk}(s_2),$$

$$g_{ijk}(s_5) \in G_{ijk}(s_5),$$

$$p_{ijk} = \hat{p}_{ijk} + g_{ijk}(s_8).$$

Any entry not listed in Table 18 is assumed to have a null effect. Likewise, in accordance with Table 15, we assume that if a property is raided, the effect is certain (i.e., all poppy is destroyed with resulting damages).

Table 18. Influence Effects on Decision Setting

Influence	Raid Probability	Poppy Income (NR)	Wheat Income	Potato Income	Tomato Income
s_2	[0.00, 0.7]	$[-0.05\hat{y}_{ijk}, 0.2\hat{y}_{ijk}]$	-	-	-
s_3	-	$[-0.15\hat{y}_{ijk}, 0.1\hat{y}_{ijk}]$	-	-	-
s_4	-	-	$[-0.25\hat{y}_{ijk}, 5.00\hat{y}_{ijk}]$	$[-0.25\hat{y}_{ijk}, 2.00\hat{y}_{ijk}]$	$[-0.25\hat{y}_{ijk}, 3.00\hat{y}_{ijk}]$
s_5	[-0.02, 0.05]	$[-0.2\hat{y}_{ijk}, 0.1\hat{y}_{ijk}]$	-	-	-

In this case study, the numerical effects of influence are considered utilizing uncertainty sets meant to capture conflicting opinions and designed to allow for alternative effects than those anticipated (e.g., s_5 decreasing the raid probability). Such conflict is historically characteristic of counternarcotic policy discussions in Afghanistan (Coll, 2019). Therefore, the notional effects in this case study are designed to be a representation of the true policy-making environment.

RDM Tailorable Components.

Having defined the uncertainty collections and sets, the next steps in the RDM approach consist of identifying a space filling design from which to build \vec{F} , a scenario generator measure, a robustness measure, an evaluation criteria, and a clustering algorithm.

From an experimental design perspective, we are considering one categorical factor with two levels (i.e., \mathcal{J}_i) and 70 additional continuous factors including each decisionmaker's CPT-parameter values, the respective influence's effects, and each crop's profitability. To accommodate this setting, we utilize a sliced Latin hypercube design because it allows for optimal space-filling properties to be pursued within and between the categorical factor levels (Ba et al., 2015). The specific design used has two slices and 1000 points in each slice, yielding a \vec{F} with 2000 futures, and was

calculated using the R-package *SLHD* with the simulated annealing parameter *iter-max* set to 20. In this manner, the resulting design matrix has favorable space-filling properties, and is convenient for calculation and adequate for illustration.

A variety of scenario generator measures could be used depending upon the precise nature of the persuader's priorities. For instance, the persuader may wish to reduce the total number of households cultivating poppy, or may wish to minimize the total number of jeribs utilized for poppy cultivation. Both fall within the larger counternarcotics objective, but are alternative metrics of success. Preferably, the selection of such a measure would be informed based upon the specific nature of the insurgency's poppy-related income stream. Absent this information, we focus on the latter scenario generator measure related to minimizing the total number of poppy cultivated jeribs. In large part, this decision is made to align with the historical priorities of US and UN policymakers (SIGAR, 2018a; UNODC, 2018a).

Furthermore, absolute regret is utilized as a robustness measure to facilitate communication of results in natural units. The interpretation of this absolute regret measure is the number of additional poppy cultivated jeribs in a given scenario from the "optimal". Moreover, the evaluation criteria selected is expected regret for similar reasons. The number of poppy cultivated jeribs in Afghanistan has a direct impact on the worldwide opium supply, and the retention of this conceptually-accessible meaning is important when communicating regret to senior policymakers.

With regard to clustering, Lempert et al. (2008) advocate for the use of two hierarchical clustering algorithms (i.e., CART and PRIM). The authors argue that each of these clustering techniques have high interpretability because they segment the area of interest into hyper-rectangular regions. Empirically, both algorithms are shown to perform well; however, CART regions are guaranteed to be disjoint whereas PRIM regions may overlap. For this reason, the CART algorithm is utilized in this

study. Moreover, to limit the depth of the classification tree, the minimum leaf size is bounded below by 35 futures in accordance with CART-application conventions.

Analysis and Results.

The result of each influence strategy on the number of jeribs produced for each future in \vec{F} is depicted in Figure 18. It can be observed that the null strategy (i.e., s_1) and most influence strategies composed of singleton actions (i.e., s_2, s_3, s_5) result in a large number of expected poppy cultivated jeribs. Among the non-singleton strategies, each of the strategies not including alternative development as an action (e.g., s_6, s_8, s_{10}, s_{13}) results in many poppy cultivated jeribs as well, indicating the importance of such actions in a successful influence strategy.




Figure 18. Distribution of Poppy Cultivated Jeribs over all Futures

The amount of variability in poppy cultivated jeribs for most of the strategies depicted in Figure 18 coincides with the historical difficulties encountered by those developing counternarcotic policy. Even strategies having the greatest potential benefit may be ineffectual when applied. For instance, if one were to assume all futures equally likely, s_{16} results in the lowest number of expected poppy cultivated jeribs even though instances exist having no decrease in cultivation relative to the null strategy.

As such, the raw distributional data on poppy cultivated areas is not sufficient to determine which of the sixteen considered strategies is robust. We cannot discern the degree of which a strategy's favorable futures coincide with that of any other. Therefore, to continue with the RDM analysis, we introduce the robustness results in Figure 19.




Figure 19. Distribution of Absolute Regret over all Futures

By observing this figure, we can discriminate better between strategies. Differences between s_{12} and s_{14} emerge based upon the sizes of their interquartile-regret ranges. A similar observation holds for s_7 and s_{16} . Notably, we also see that s_7 , s_{12} , s_{14} , and s_{16} all have an expected absolute regret of zero. However, s_{14} and s_{16} have the smallest interquartile-regret ranges and are leading contenders for the candidate strategy for policy implementation. Of the two, s_{16} has the lower 90th percentile regret and, for this reason, is chosen as s_{cand} .

With a candidate strategy selected, the RDM methodology calls for hedging options to be explored. As in Lempert et al. (2006), each future's output is transformed into a binary response. The binary response coincides with an absolute regret threshold set forth by senior policymakers. We assume for this use case that such a limit is 2500 jeribs in regret. That is, if the number of poppy cultivated jeribs for s_{cand} exceeds 2500 jeribs of the best strategy for some future, then that future is considered

a vulnerability of s_{cand} .

Utilizing this technique, the CART algorithm clusters the space into eight regions of interest. These regions are separated based on the effectiveness of select counternarcotics actions, and a subset of demographics' perceptions on crop income. However, one cluster stands out as having a high density of vulnerabilities and high average absolute regret. It is characterized by the simultaneous effectiveness of interdiction and political support mobilization on poppy income estimates, the ineffectiveness of eradication messaging on the perceived raid probabilities, the DVI demographic believing poppy and wheat income high in ample rain conditions conditions, and the RF demographic regarding poppy income as high in ample rain conditions and tomato income as high in drought conditions. With this vulnerable narrative identified, the tradeoffs between utilizing s_{16} or some other strategy in terms of expected absolute regret can be examined.




Figure 20. Cluster Tradeoffs Between s_{16} and Other Strategies

Figure 20 illustrates the tradeoffs between utilizing s_{16} vis-à-vis the alternatives in terms of expected absolute regret. The coordinate location of a strategy is determined by the expected absolute regret via two narratives: the vulnerability cluster and the remaining futures. The four strategies not depicted in Figure 20 (i.e., s_1, s_3, s_5, s_{10}) exceed 3×10^5 in both narratives. A frontier of non-dominated strategy performance is formed by s_{12} and s_{16} , as depicted by the solid line. This indicates that only s_{12} and s_{16} are non-dominated strategies under this narrative bifurcation. It can be observed that a persuader's preference between the two depends upon the respective weight (perhaps in terms of probability) placed on the vulnerable scenario and its complement. If weighted equally, a strategy on the dotted line closest to the origin is preferred. However, as one narrative is given more weight, the preference shifts to strategies close to the origin on one side of the line. Therefore, the more (or less) weight the persuader respectively places on the vulnerable scenario, the more (less) attractive strategy s_{12} becomes vis-à-vis strategy s_{16} .

At this point, the RDM process can be repeated with s_{12} as s_{cand} and, potentially, with the inclusion of additional information (e.g., more crops or additional influence actions). Conversely, the RDM analysis can terminate with a choice between s_{12} and s_{16} based upon the persuader's beliefs on the vulnerable and the remaining future narratives for expected absolute regret.

Moreover, the RDM procedure could also be repeated utilizing a distinct scenario generator metric. As mentioned in Section 4.3, a variety of counternarcotics measures exist to assess a strategy's success. Besides the total number of poppy-cultivated jeribs, the total kilograms of opium produced from a harvest has historically been a relevant measure to policy-makers. Therefore, utilizing historical kilograms per hectare yield data (UNODC, 2018b), estimated best- and worst-case empirical CDFs (e.g., Figure 21) for the tons of opium from the Badakhshan poppy harvest could

inform an alternative RDM analysis under a different measure defined by a senior policy-maker.

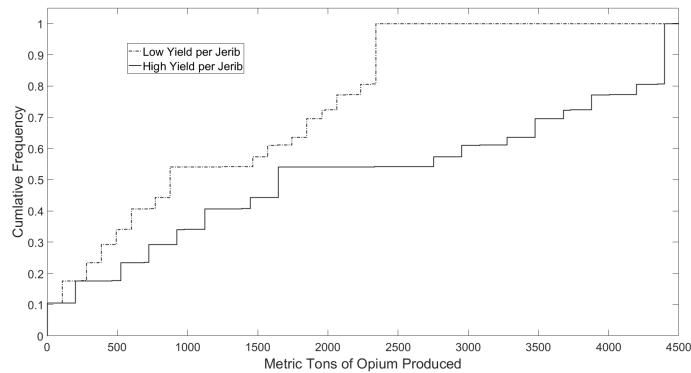


Figure 21. Empirical CDFs for s_{12} with Best- and Worst-Case Opium Yields

Discussion and Limitations.

The selection of s_{16} as an initial candidate strategy is reinforced by counternarcotic doctrine proposed by Ward and Byrd (2004) and SIGAR (2018a) who envision it as a multifaceted effort. Furthermore, the choice between s_{12} and s_{16} coincides with a core issues at the heart of the Afghan conflict: the effectiveness of interdiction and security cooperation efforts (SIGAR, 2018a). In our use case, if a senior policy-maker believes interdiction efforts can decrease farmers' poppy net revenues and, they are also optimistic regarding the capability of Afghan partners and the coalition's ability to support them, then action s_{12} should be adopted. Otherwise, s_{16} is the preferred strategy.

The modeling methodology also addresses some criticisms by Mansfield and Fishstein (2016) pertaining to faulty assumptions in Afghan counternarcotic policy development. That is, we do not treat the Afghan farmer as a homogeneous entity maximizing gross profit, nor one who makes a binary decision between opium poppy and wheat. Instead our approach addresses the diversity of Afghanistan demographics, considers net profits, and examines a variety of crop selection profiles.

However, the modeling methodology utilized is not a perfect representation of reality. Although we addressed the diversity of Badakhshan farmers, a higher fidelity characterization may be required to inform policy decisions. Badakhshan is a multiethnic region comprised of Tajiks, Uzbeks, and Pashtos, among others (Fishstein, 2014). The effect of a farmer's demographic on a crop decision may be non-trivial, especially when concerned with emotional stimuli (Gladwell, 2008). Moreover, a richer characterization of the crops available for planting may also be required using district rather than provincial level data. A diverse set of crop options are available to Afghan farmers, and deterring poppy cultivation may require creative incentives among them. Unfortunately, the data required to inform such decisions at this granular level is not available in open source collections and would require an expansive data gathering effort analogous to other UNODC projects.

Our models also only consider a portion of the interconnected Afghan economy, of which farming is a critical component. It provides land-access (via sharecropping) and income (via day labor) to the rural poor. As such, the crop selection of Afghan landowners can have far reaching effects. Such effects were observed during the implementation of the HFZ. While the program resulted in landowners deciding to plant less poppy, these decisions resulted in less profitable sharecropping agreements and less demand for on-farm labor (Mansfield, 2017).

As such, the program coincided with a migration of the landless poor into former desert areas (e.g., Bakwa in Farah province) and an increase in poppy cultivation in these regions (SIGAR, 2018a). Just as such effects were unforeseen by the planners of the HFZ, so too would they be unforeseen in the analysis presented in this research. The decisions of the landless poor could be included and linked to the our analysis but, ultimately, the scale of the influence model must be determined based upon the persuader's underlying objectives.

The influence models presented herein are static, closed system approximations of a component in an open, dynamic system. The entirety of the dynamic system cannot be modeled. Instead, the influence model should be contained to represent only the relevant elements of the full system. In this study, we assume such elements to be restricted to the landowners themselves.

Finally, the Badakhshan landowner influence model is designed to affect change in a single growing season. Such a time frame is far too short upon which to build a comprehensive counternarcotics policy. Historically, successful policies have taken decades versus months or years (Ward and Byrd, 2004) and, in this time frame, external market factors not incorporated in this research become increasingly relevant. Therefore, the models illustrated herein do not yield a panacea to stop poppy cultivation, but they facilitate the development of counternarcotic policy toward incremental, gradual change.

4.4 Conclusions

In this research, we presented a framework for modeling influence under parametric and structural uncertainty and illustrated how robust decisions in this setting can be identified. In this manner, a persuader can generate a robust *nudge* strategy (Thaler and Sunstein, 2009) that relies upon System 1 (i.e., automatic thought), System 2 (i.e., reflective thought) influence, or a mixture thereof (Kahneman, 2011).

The utility of our methodology was demonstrated on the global, contemporary problem posed by the illicit opium economy. Specific emphasis was applied to reducing opium poppy supply from the world's leading poppy producing nation by targeting the Afghan province of Badakhshan. This province is an appealing target for counternarcotic efforts because, in accordance with SIGAR (2018a) guidance, it is relatively well-controlled by the national government. However, insurgent groups

have a large enough presence regionally that they may be generating significant income from the province's illicit activities (Chughtai, 2018).

In addition to providing an alternative model for counternarcotic policy development, the Badakhshan use case illustrates how an influence model can be developed under uncertainty. The discussion in Section 4.3 exemplifies the level of empathy and situational awareness required both to infer the underlying decision setting and to estimate the effects of influence upon it. From a technical perspective, the use case also demonstrates how the RDM framework can be tailored to fit a particular influence problem and how iterative rounds of analysis ought to be performed.

However, it is also the technical application of RDM for which much work remains to refine the methodology. The development of \vec{F} is based upon a space-filling design; however, there are a variety of options from which a modeler can choose. Future inquiry is required to determine when one space-filling design is preferable to another and what sampling density is suitable. Another open question pertains to the usefulness of optimal space-filling designs in this setting. The scenario generator is not particularly demanding from a computational perspective; in fact, it is a relatively direct application of CPT. Therefore, future inquiry into the tradeoff between computational effort and solution quality via random sampling, sub-optimal space-filling designs, or optimal space-filling designs is a promising endeavor.

V. Leveraging Behavioral Game Theory to Inform Military Operations Planning

Abstract

Since Thomas Schelling published *The Strategy of Conflict* (1960), the study of game theory and international relations have been closely linked. Developments in the former often trigger analytical changes in the latter, as evidenced by the recent behavioral and psychological focus among some international relations and defense economics scholars. Despite this connection, decisions regarding military operations have rarely been influenced by game theoretic analysis, a fact often attributed to standard game theory's normative nature. Therefore, this research applies selected behavioral game theoretic solution techniques to classical interstate conflict games, demonstrating their utility to inform the planning of military operations. By reexamining classic Cold War deterrence models and other interstate conflict games, we demonstrate how modern game theoretic techniques based upon agent psychology, as well as the ability of agents to think strategically or learn from past experience, can provide additional insights beyond what can be derived via perfect rationality analysis. These demonstrations illustrate how behaviorally focused methods can incorporate the uncertainty related to human decisionmakers into analysis and highlight the alternative insights a bounded rationality approach can generate for military operations planning.

5.1 Introduction

For nearly two weeks at the end of October 1962, the events surrounding the Cuban Missile Crisis arguably constituted the single most consequential decision setting in history. Fortunately, Soviet premier Nikita S. Khrushchev and American president

John F. Kennedy both chose to back away from decisions that would escalate the conflict, possibly into a nuclear armageddon. A disaster was averted.

The characteristics of a leader determine, at least in part, their strategic decisions. History has shown that a leader's intellect, temperament, background, biases, and other psychological factors play a critical role throughout the course of a conflict and in its resolution. Therefore, more than fifty years after the Cuban Missile Crisis, one cannot help but wonder... what if Joseph Stalin had not passed away? What if Richard Nixon had won the 1960 American presidential election? What if the *Bay of Pigs* had not eroded President Kennedy's trust in his military advisors? How would the successive decisions made by these world leaders have differed? Unfortunately, traditional game theoretic concepts used to analyze such situations provide no insight to answer such questions.

In formal models of interstate conflict, states are often modeled as strategic-thinking entities attempting to maximize their self-interest (Schelling, 1960, 1966; Powell, 1990; Zagare and Kilgour, 2000). Insights are historically gained through an application of some traditional game theoretic equilibrium concept. However, the concept of an equilibrium imposes an additional assumption. Players are assumed to form their responses based upon accurate beliefs of what their adversaries actually do; that is, all players are *mutually consistent* (Camerer et al., 2004). An equivalent, and more formal exploration of these conditions is considered by Aumann and Brandenburger (1995). Although alternative explanations exist to describe how these conditions arise (e.g., folk wisdom or instinct), one valid interpretation is to consider decisionmakers in these models as perfectly rational actors.

Under such an interpretation, the human aspects of the respective nations' leaders are not considered. People are generally overconfident (Kahneman, 2011; Johnson and Fowler, 2011). Their decisionmaking is affected by emotional factors (Thaler,

1999; Shiv et al., 2005), and they are unable to objectively perceive uncertainty (Kahneman and Tversky, 1972, 1979, 1982; Kontek and Lewandowski, 2017). The interpretation of humans (e.g., national leaders) as perfectly rational decisionmakers is further complicated by the prevalence of mental health disorders in society. Wittchen and Jacobi (2011) estimated that, within any contemporary twelve-month period, 27% of adults in the European Union suffered from a mental disorder. Davidson et al. (2006) estimated that 49% of American presidents between 1776 and 1974 suffered from some psychiatric disorder (e.g., depression, alcoholism), and McDermott (2007) illustrated the profound effect medical and psychological illness have exerted on presidential decisionmaking.

For these reasons, recent research in international relations and defense economics has begun to consider behavioral and psychological factors (e.g. see Pittel and Rubbelke, 2012; Phillips and Pohl, 2017; Horowitz and Fuhrmann, 2018; Apolte, 2019). Rathbun et al. (2016) and Kertzer et al. (2018) respectively studied behavioral effects on foreign-policy beliefs and costly signaling in international relations. Little and Zeitzoff (2017) incorporated behavioral factors into bargaining models via evolutionary preferences. Hafner-Burton et al. (2017) provided an extensive review of relevant behavioral theories, previous studies of their implementation, and their collective implication for the study of international relations. Other recent examples of behavioral studies questioning the role of rationality in political decisions include the work of Caplan (2006), Andonie and Diermeier (2017), Thomas (2018), Lee and Clark (2018), and Tóth and Chytilek (2018).

However, such a behavioral paradigm shift has not yet manifested itself in game theoretic applications to military operations planning. The lack of emphasis on players' humanity in traditional solution concepts has discouraged the application of game theory to the planning of military campaigns. "Homo economicus rather than deontic

logic have deterred [war]gamers from studying game theory, and thus the perceived value of applying game theory to [war]gaming has been limited” (Hanley Jr., 2017b). Whereas many scholars in recent decades have seen the potential of incorporating game theory into the process (e.g. see Leibowitz and Lieberman, 1960; Taylor, 1978; Athans, 1987; Cruz et al., 2001; McEneaney et al., 2004; Boardman et al., 2017), actually doing so has been limited by the military desire to “anticipate how the key players may act” (Hanley Jr., 2017b) and the difficulty of anticipating human behavior when assuming perfectly rationality of actors. Conversely, behavioral game theory (BGT) “expands analytical theory by adding emotion, mistakes, limited foresight, doubts of how smart others are, and learning to analytical game theory” (Camerer, 2011) and focuses on describing how players actually interact.

Of the recent BGT developments, two frameworks are of particular importance for military operations planning: the Cognitive Hierarchy (CH) model and Experience Weighted Attraction (EWA) model. Camerer et al. (2004) developed the CH model that formalized the intuition of Selten (1998) that the “natural way of looking at game situations... is [through] a step-by-step reasoning procedure”. With relatively few assumptions, the authors develop an empirically accurate model of play for one-shot games based upon a parameter characterizing the players’ strategic thinking abilities. By comparison, the EWA model developed by Camerer and Ho (1998, 1999) describes the empirical learning behavior of agents in a variety of repeated games. The EWA model is defined with psychological parameters identifying the players’ tendencies toward regret, forgetfulness, and myopia.

From a military operations perspective, these models are particularly useful when planning for a one-shot engagement, a new type of hostility, or a recurring conflict. Instead of removing the underlying human uncertainties, these models incorporate them into an agent-based approach, thereby addressing two deficiencies associated

with legacy analysis techniques (Hanley Jr., 2017a). As such, in this research, we illustrate how CH and, when appropriate, EWA can be effectively utilized to garner insights for the conduct of military operations.

The remainder of this manuscript is structured as follows. We begin in Section 5.2 by reviewing each behavioral game theory model and their respective parameters. These models are applied in Section 5.3 to classic interstate conflict games, illustrating the alternative insights they derive relative to traditional methods and, in doing so, exhibiting their utility for military operations planning in an accessible manner. We also discuss how behavioral uncertainty can be incorporated to inform the planning of a military operation and balance between the two command styles discussed by Haywood (1954): planning against enemy capability vis-à-vis planning against enemy intent (i.e., what the enemy *is able to* do versus what the enemy *will* do). Finally, we discuss additional BGT models that can be used to inform military operations in Section 5.4 and provide concluding remarks in Section 5.5.

5.2 Behavioral Game Theory

“Behavioral Game Theory stands alone in blending experimental evidence and psychology in a mathematical theory of normal strategic behavior” (Camerer, 2011). The field takes an experimental economic approach to game theory, and it utilizes empirical results to guide theory. In this section, we review two prominent models in the discipline: one which characterizes behavior in one-shot, normal form games, and the other describing learning in repeated play. These models provide a basis of play rooted in empirical results, but also provide more definitive predictions than their standard equilibrium counterparts. For a given τ -parameter characterizing the distribution of opponents’ depths of strategic thought, the CH model provides a single estimate of play, and when simulated many times, the EWA model provides a range

of expected behavior. Likewise, these behavioral methods are readily adaptable to games having a large number of players, whereas finding equilibrium in such situations is a computationally burdensome task (Shoham and Leyton-Brown, 2008).

Cognitive Hierarchy.

From the earliest application of game theory to interstate conflict, there has existed the concept of step-by-step reasoning. Schelling (1960) referred to this general structure with regard to the “reciprocal fear of surprise attack”, describing how the Soviet Union and the United States interact in a game of preemptive attack. Schelling described the problem as a process produced by successive cycles of “He thinks we think he thinks we think... he thinks we think he’ll attack; so he thinks we shall; so he will; so we must”. Schelling (1960) noted that this structure can be formalized using a potentially infinite set of probabilities composed of the products of the following event probabilities and their reciprocals: P_1 (i.e., the probability the opponent prefers to attack), P_2 (i.e., the probability opponent thinks we prefer to attack), P_3 (i.e., the probability opponent thinks I believe they prefer to attack), and so on.

In Schelling’s estimation, “the trouble with the formulation is that nothing generates the series” and that “each probability is an ad hoc estimate”. Instead, he sought a reformulation of the problem to gain insights. However, the Cognitive Hierarchy model set forth by Camerer et al. (2004) allows for analysis of player interaction in its original step-by-step reasoning form.

The CH model assumes that players are defined by the number of reasoning steps they compute in selecting an action, and their beliefs on the number of steps computed by other players. Agents are overconfident, and they do not realize that other agents can use as many reasoning steps as they do (i.e., a k -step player believes all others are $[k-1]$ -step players or less). A game is assumed to be characterized by a true probability

distribution over the number of reasoning steps a player utilizes. Generally, a Poisson distribution defined by the parameter τ is used. The larger the value of τ , the more reasoning steps players utilize, on average. As players are overconfident, they are unable to perceive this true distribution. Instead, Camerer et al. (2004) assumed the subjective belief a k -step player assigns to encountering an h -step player is equal to

$$g_k(h) = \frac{f(h)}{\sum_{l=0}^{k-1} f(l)} \quad (36)$$

where $f(h)$ is the true probability of an h -step player being encountered from the underlying Poisson distribution. That is, the true probabilities are normalized to form a k -step player's beliefs. Utilizing these beliefs, a player's strategy is calculated by maximizing their expected payoff. Namely, a strategy j for player i is selected such that s_i^j maximizes

$$E(\pi(s_i^j)) = \sum_{j'=1}^{m-i} \pi(s_i^j, s_{-i}^{j'}) \left[\sum_{h=0}^{k-1} g_k(h) P_h(s_{-i}^{j'}) \right] \quad (37)$$

where $E(\pi(s_i^j))$ is player i 's expected payoff for playing strategy j ; $\pi(s_i^j, s_{-i}^{j'})$ is player i 's payoff for playing strategy j when his opponents play the strategy vector j' ; $m-i$ is the number of possible opponent strategy vectors; and $P_h(s_{-i}^{j'})$ is the probability that h -step players would actual commit to strategy j' . If multiple strategies are found to achieve the same maximum payoff, the player is assumed to randomize equally among them.

In this way, the CH model formalizes the step-by-step reasoning process by limiting its depth and forcing convergence as k increases. As k grows sufficiently large, a k -step and a $[k+1]$ -step player have the same beliefs and therefore take the same action(s). In practice, a game is solved via CH recursively by starting with 0-step players and iterating up to this k -value. Thus, the assumption of 0-step player behavior is the

bedrock for all higher order player actions. Such players are assumed to not think strategically and randomize over the strategy space. In keeping with Camerer et al. (2004), we assume 0-step players randomize over this space uniformly.

The CH model primarily characterizes behavior in a game by the Poisson parameter τ . In their study of a variety of games, including the beauty contest and market entry games, Camerer et al. (2004) found that common estimates of τ are approximately in the range $[1, 2]$ with the value 1.5 reliably predicting expected behavior in new games. Likewise, as $\tau \rightarrow \infty$ the CH model will converge to any Nash equilibrium reached by finitely deleting weakly dominated strategies but, in general, such convergence is not guaranteed.

In a military operations planning setting, it may prove difficult to derive enough empirical evidence to estimate τ directly. However, qualitative estimates can be identified by heuristically assessing the level of strategic thinking utilized by the interacting entities. Alternatively, the interaction can be readily examined under a large number of τ -estimates to obtain a range of expected behavior. Such an approach is illustrated in Section 5.3 by examining all τ in the set $\{0, 0.1, \dots, 50\}$.

Furthermore, the CH model is fit with τ -values corresponding to a population of individual decisionmakers. In a military operations setting, the ultimate decision authority may reside in an individual, but input is invariably received from a staff of advisors. The decision process itself is often cooperative in nature. Therefore, it is of interest for military operations planners to understand how this dynamic affects the underlying τ -value in one-shot games. We hypothesize that the collective input from advisors will serve to increase the depth of strategic reasoning, but experiments must be conducted to support or refute this conjecture.

Experience Weighted Attraction.

Camerer (2011) noted that “there are no interesting games in which subjects reach a predicted equilibrium immediately. And there are no games so complicated that subjects do not converge in the direction of equilibrium”. However, standard game theory does not explicitly address the expected behavior of boundedly rational agents, nor does it address the temporal aspect of equilibration. In some games, players quickly adapt and move to an equilibrium state but, in other games, it is a process that may require years or decades (Camerer, 2011).

The EWA model developed by Camerer and Ho (1999) attempts to address these concerns. It is a learning model that can be tuned to characterize the path players take towards an equilibrium profile over repeated play. The EWA model assumes players use reason to adopt a starting strategy profile which then evolves due to both beliefs formed and reinforcement received over repeated interaction.

The EWA model is centered upon two key concepts: (1) an accumulated experience variable, and (2) attraction values toward given strategies. At time zero, each player i is assumed to have some initial experience value, $N_i(0)$, and some starting attractions toward each strategy j , $A_i^j(0)$. The initial experience value is assumed to be derived from experience playing different games and introspection, whereas the initial attractions are assumed to come from some strategic reasoning process (e.g., the CH model). Subsequent experience variables for player i at time t are represented as

$$N_i(t) = \phi_i(1 - \kappa_i)N_i(t - 1) + 1, \quad t \geq 1. \quad (38)$$

wherein ϕ_i is a decay factor representing loss of salience of previous interactions, and κ_i is a factor describing the level at which the strategy space is explored or exploited. That is, κ_i -values closer to 0 represent a tendency to explore, whereas κ_i -values closer

to 1 represent a more exploitive framework. Collectively, the product of these two parameters describes how past experience is discounted.

The attraction value player i allots to a given strategy j at time t utilizes the experience variable and is calculated as follows:

$$A_i^j(t) = \frac{\phi_i N_i(t-1) A_i^j(t-1) + [\delta_i + (1 - \delta_i) I(s_i^j, s_i(t))] \pi_i(s_i^j, s_{-i}(t))}{N(t)} \quad (39)$$

such that δ_i is an imagination factor placed on foregone payoffs, s_i^j is the strategy of corresponding to $A_i^j(t)$ support, $s_i(t)$ is the strategy actually played by player i in period t , and $I(s_i^j, s_i(t))$ is an indicator function turning on and off the imagination effect, as appropriate. That is, $I(s_i^j, s_i(t)) = 1$ if $s_i^j = s_i(t)$, and 0 otherwise.

These attraction values are then used to update the probabilities with which each player utilizes a given strategy by

$$P_i^j(t+1) = \frac{e^{\lambda A_i^j(t)}}{\sum_{k=1}^{m_i} e^{\lambda A_i^k(t)}} \quad (40)$$

wherein λ represents sensitivity to attractions and may be influenced by perception, rate of computational errors, or unobserved payoff components (e.g., desire for variety of play).

The EWA model has been shown to generalize a variety of other learning algorithms (e.g., reinforcement learning, belief learning, and weighted fictitious play), depending on the values of κ_i , ϕ_i , and δ_i . If these requisite parameters are known (or assumed), learning behavior is simulated by successively generating player strategies based on the values of $P_i^j(t)$ and updating $N_i(t)$, $A_i^j(t)$, and $P_i^j(t+1)$. Expected insights from the model can then be gained via the aggregation of outcomes over multiple simulation runs.

Alternatively, if observed behavior is known, the requisite parameters may be fit by minimizing the likelihood function in a process analogous to that demonstrated by Camerer and Ho (1998), or by performing multiple runs with a space-filling design and selecting the parameter tuple minimizing the mean squared error of averaged simulated path, similar to Roth and Erev (1998). The former is more computationally efficient, but the latter is easier to implement.

With regard to the use of EWA in military operations planning, the characterization of learning behavior is extremely useful because many problems in combat do not occur in a one-shot setting but repeat multiple times. Whereas an invasion or an armistice may occur in a one-shot setting, other activities (e.g., reconnaissance) are frequently revisited. For such missions, learning plays a definitive role in the interaction. Utilization of the EWA model or one of its extensions (e.g. Camerer et al., 2002; Ho et al., 2007) to the planning of military operations can incorporate this element of repeated play and may provide heretofore hidden insights.

5.3 Applications to Military Operations Planning

The military operations planning process can be characterized by the following five steps: (1) receipt of the mission, (2) development of situational awareness and available courses of action, (3) analysis of opposing courses of action, (4) comparison of courses of action, (5) selection of a course of action. Although the delineations between the aforementioned activities may differ, such a general framework has been used in the United States since before World War II (Haywood, 1954) and is a hallmark of modern, Western operations planning (Australian Department of Defence, 2009; U.K. Ministry of Defense, 2013; SHAPE, 2013; U.S. Joint Chiefs of Staff, 2017b).

Given such commonalities, we illustrate in this section how BGT can help analyze and compare courses of actions for a military operation. Generally speaking, we

refrain from adapting the lexicon of any particular planning paradigm (e.g., the Joint Operations Planning Process); however, we do find it useful to highlight how each of the examples relate to the six-phase Continuum of Military Operations: Phase 0 (Shaping the Environment), Phase 1 (Deterring the Enemy), Phase 2 (Seizing the Initiative), Phase 3 (Dominating the Enemy), Phase 4 (Stabilizing the Environment), and Phase 5 (Enabling Civil Authority). For further information regarding these phases, we refer an interested reader to *Joint Publication 3-0: Joint Operations* (U.S. Joint Chiefs of Staff, 2017a).

In Sections 5.3-5.3, we treat the military planners as endogenous entities in the respective BGT frameworks. That is, planning activities are conducted with the incorporation of friendly and adversarial bounded rationality. Such a perspective is of particular utility in decentralized execution structures because a higher headquarters does not always seek to directly control subordinate commanders' decisions; instead, these commanders are often granted a degree of autonomy to account for the velocity of the decision making environment (Shamir, 2010). In these sections, we also illustrate how planners can incorporate uncertainty regarding behavioral parameters (i.e., behavioral uncertainty) into their analyses.

Conversely, in Section 5.3, we consider a setting wherein the military planners are exogenous of the BGT frameworks. These planners are aware of the adversary's bounded rationality and seek to exploit it to their benefit. Such an approach is most appropriate for centralized command and execution. For these settings, we illustrate how planners can incorporate behavioral uncertainty to select robust courses of actions in view of the enemy's capabilities and intent.

Each of these sections utilize BGT in the analysis of well-known, interstate conflict models. Although these models are classical rather than contemporary, we utilize them to ensure the accessibility of our results to a broad readership by striking a bal-

ance between military relevance and technicality of illustration. However, in Section 5.4, we discuss how the BGT methods utilized herein and others introduced in Section 5.4 can be utilized in more complex game theoretic settings.

Brinkmanship: Nuclear Crisis Game.

We begin by analyzing a well-known metaphor for brinkmanship, namely the the two-player, normal form game of Chicken in a one-shot setting (Russell, 1959; Kahn, 1960; Rapoport and Chammah, 1966). Although often framed in the context of nuclear deterrence, a variety of military activities at the tactical level exhibit a similar structure (e.g., confrontations during Freedom of Navigation operations). As such, we take the perspective of a higher headquarters planning for such a deterrent activity in Phase 1 of the Continuum of Military Operations. The decision authority for this interaction has been delegated to a subordinate commander, and the higher headquarters wishes to understand the potential behavior of friendly and enemy forces arising from a particular course of action.

The specific payoffs we utilize for this analysis can be found in Table 19. Should this matrix game be analyzed with the assumption of perfect rationality, there are two pure Nash equilibria and one mixed Nash equilibrium: (1) the row player holds firm and the column player deescalates, (2) the row player deescalates and the column player holds firm, or (3) both players deescalate with probability 0.9, and hold firm with probability 0.1.

Table 19. Two-player Nuclear Crisis Game

	Deescalate	Hold Firm
Deescalate	(0,0)	(-1,1)
Hold Firm	(1,-1)	(-10,-10)

If this assumption is relaxed, and players are assumed to be boundedly rational by utilizing the step-by-step reasoning of the CH model, we arrive at similar, albeit differing results. Since the higher headquarters is uncertain of the strategic thinking ability of both forces' leaders, the CH model is utilized with all instantiations of τ in the set $\{0, 0.1, \dots, 50\}$. Figure 22 displays the probability that a player will choose to deescalate as a function of τ ; the probability of holding firm is its compliment.

The three horizontal lines on the graph represent the three Nash equilibria. The symmetry of the game lends itself to identical strategy profiles under CH for the row and column players. Thus, Figure 22 applies to both players. Utilizing the suggested value of $\tau = 1.5$ for untested games (Camerer et al., 2004), the CH model predicts that a player will deescalate with probability 0.8884. If this value of τ is correct, then we should expect the mixed form equilibrium to predict behavior well. Furthermore, the CH model has enabled differentiation between the three Nash equilibria. Our results suggest that when $\tau \rightarrow \infty$ the players' strategies converge to the mixed Nash equilibrium.




Figure 22. Player strategies as a function of τ

However, the CH model also yields an unexpected result. We can assume that war occurs in this game if both players simultaneously choose to hold firm. Therefore, the probability of war is the associated product of each players' hold firm probabilities. As the average depth of strategic thinking (i.e., τ) increases, war becomes improbable,

converging to a 1% chance. However, there are oscillatory patterns which make marginal increases in strategic thought dangerous. For instance, with $\tau = 1.5$ there is a 1.2% chance of war; however, a relatively marginal increase to $\tau = 2$ corresponds to a 6.2% chance of war. This illustrates a counterintuitive instance wherein a more strategic population of decisionmakers is more likely to take action resulting in a worse collective outcome. The importance of such results to military operations planning is apparent when considering that this example is reminiscent of how Tuchman (1962) described the onset of World War I.




Figure 23. Probability of war as a function of τ

We now focus on the subject of learning and consider how player strategies might evolve when two forces repeatedly engage in brinkmanship. The EWA model is designed to fit observed data. However, if no data exist, a strategist may be able to qualitatively assess each actor's tendency toward regret (δ_i), forgetfulness (ϕ_i), and myopia (κ_i). Alternatively, they could utilize regressed parameters from other games if they believe behavior will be similar. Therefore, in this notional example, we utilize the parameters fit by Camerer et al. (2002) on the continental divide game such that, for any player i , we have $\phi_i = 0.61$, $\kappa_i = 1$, and $\delta_i = 0.75$. The initial attractions are found by inserting the probabilities derived from the CH model with $\tau = 1.5$ into equation (40) and solving for $A_i^j(0)$. Likewise, for simplicity, we assume that for each player i we have $\lambda_i = 1$ and $N_i(0) = 1$. One hundred simulations were run for one hundred interactions each to explore how players of these types might learn in this

game. The results of these simulations can be seen in Figure 24 .

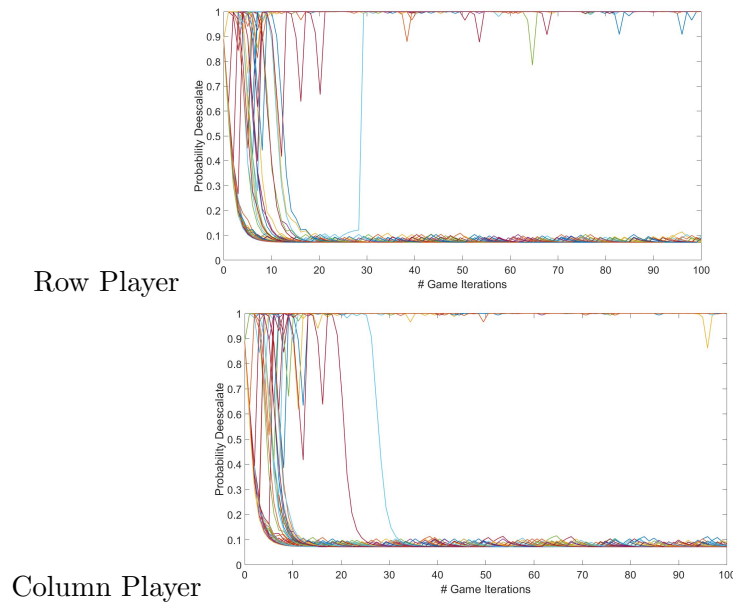


Figure 24. Players strategies over time in 100 simulated EWA runs

The plots within Figure 3 illustrate a set of players that diverge from the mixed strategy equilibrium and converge in the direction of one of the pure strategy equilibriums. That is, as one player decreases the probability of holding firm, the other increases their probability of holding firm, and vice versa. Of note, the lowest the probability of holding firm reached in all simulations was 0.07. While it never reached zero, this is not an artifact of the EWA model but of the parameters chosen. For instance, if the ϕ_i -values are increased to 0.81, players reach a zero probability of holding firm.

The results depicted in Figure 24 are interesting, especially because they show some instances wherein the players attempt to switch between the true pure strategy equilibria. However, they also implicitly assume the brinkmanship game is repeatedly played even if both sides hold firm. Alternatively, if one assumes that war occurs if both sides decide to hold firm simultaneously, the histogram in Figure 25 illustrates the number of time periods our brinkmanship game was played until either a war

occurred or the 100 limit maximum was reached, in bins of 5 time periods. It can be observed that the majority of simulations end with a positive outcome: no war occurs. However, approximately 30% of the simulations end prior to 10 stage game iterations.




Figure 25. Number of time periods before nuclear exchange

The results of this BGT analysis, in effect, categorize the Nash equilibria. When the game is first played, the CH model predicts that strategic reasoning enables players to arrive at the mixed Nash equilibrium strategy. As agents repeatedly play the game, strategies diverge from this profile and converge toward one of the pure Nash equilibria wherein one player deescalates with certainty. In general, this convergence process takes time but is complete after 50 iterations of play. This result implies that the risk of war is at its highest in early stages of brinkmanship and reduces drastically if war is not triggered at this early juncture. Such an observation is supported by Figure 25 wherein war occurs infrequently after 10 repetitions of the stage game. Therefore, for this game, the CH and EWA models do not necessarily supplant analysis from a perfect rationality perspective but augment it.

Target Selection: D-Day Game.

We again consider a higher headquarters examining the ramifications of a particular course of action on the interaction between a subordinate commander and

the enemy. Assuming an interaction occurring during Phase 3 of the Continuum of Military Operations, a subordinate commander is tasked to achieve some objective and is delegated targeting responsibilities. The enemy wishes to limit the damage of such aggression by hardening the targeted location.

Such a setting can be modeled in a manner analogous to a simple Colonel Blotto game. For example, the Allied invasion of German-occupied France in WWII can be modeled via a modified Matching Pennies game (Crawford, 2003). The Allies and the Germans can choose to attack or heavily defend, respectively, one of three sites: Calais, Normandy, or Brittany. If the Allies invade a site the Germans have chosen to heavily defend, they lose. Otherwise, they win. However, the game differs from the standard Matching Pennies games in that there are three strategies, it is not zero-sum, and selection of some strategies may incur an additional cost. The associated payoff matrix, adapted from Kydd (2015), can found in Table 20.

Table 20. D-Day Game

		Germans		
		Calais	Normandy	Brittany
Allies	Calais	(0,1)	(1,0)	(1,0)
	Normandy	(1- c_N ,0)	(- c_N ,1)	(1- c_N ,0)
	Brittany	(1- c_B ,0)	(1- c_B ,0)	(- c_B ,1)

For illustration, we assume that $c_N = 0.25$ and $c_B = 0.4$. This instance of the D-Day game does not have a pure strategy (strict) Nash equilibrium, but it has a mixed equilibrium with $p(\text{Calais}) = 0.55$, $p(\text{Normandy}) = 0.3$, and $p(\text{Brittany}) = 0.15$ for the Germans, and the Allies mixing uniformly over their available strategies. Under this equilibrium profile, the Allies attack a well-fortified location (i.e., attack a position heavily defended by the Germans) with probability 0.334.

By its nature, the D-Day game can only be played once. As such, we utilize the CH model to understand how different levels of strategic thought affect the game's

likely outcome, and compare these results with analysis under the assumption of perfect rationality.

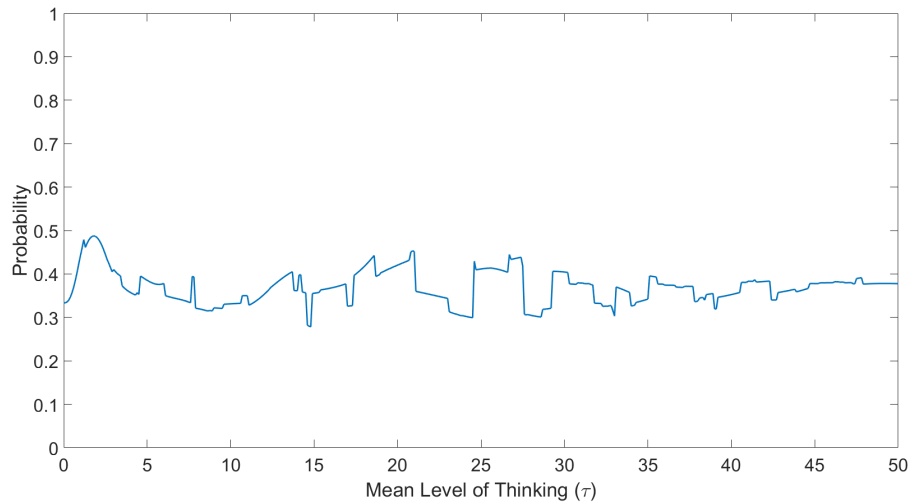


Figure 26. Probability Allies attack a well-fortified location as a function of τ

We begin by analyzing the probability that the Allies attack a well-fortified location. This probability as a function of τ can be seen in Figure 26. With smaller τ -values, this probability appears to vary erratically, but begins to behave more stably as τ increases. However, the probability of attacking a well-fortified location in the CH model appears to converge to approximately 0.38 as opposed to the 0.334 in the Nash equilibrium, indicating the CH model is yielding a different strategy. The resulting CH strategies of each player for $\tau \in \{0, 0.1, \dots, 50\}$ can be seen in Figures 27 and 28.

The CH model provides different insights than would be expected from an analysis assuming perfect rationality. At 0-level thinking, the Germans completely randomize across their strategies and as they increase in strategic thinking (along with the Allies), their play begins to diverge from this profile. However, their strategy does not converge to the mixed equilibrium. In fact, at $\tau = 50$, the Germans are expected to randomize as follows: $p(\text{Calais}) = 0.4053$, $p(\text{Normandy}) = 0.5296$, and $p(\text{Brittany}) = 0.0651$. In turn, the Allies, who begin at their equilibrium strategy,




Figure 27. Allies' strategies with $\tau = 0, 0.1, 0.2, \dots, 50$

largely avoid this profile for other levels of τ . For example, at $\tau = 50$, the Allies are expected to randomize with $p(\text{Calais}) = 0.3966$, $p(\text{Normandy}) = 0.3824$, and $p(\text{Brittany}) = 0.2210$.

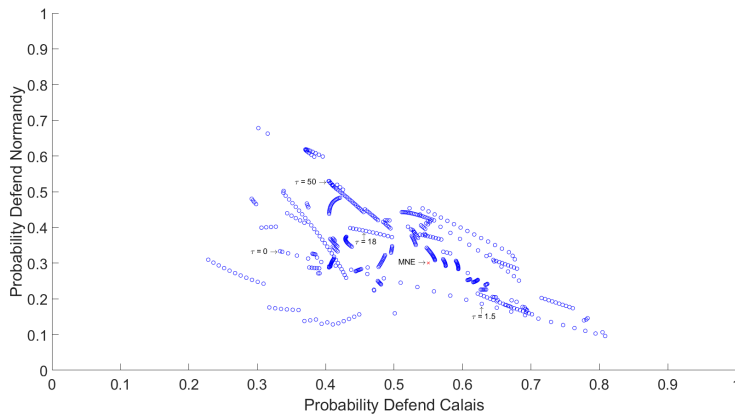


Figure 28. Germans' strategies with $\tau = 0, 0.1, 0.2, \dots, 50$

Therefore, unlike the Nuclear Crisis game, the CH model does not appear to converge towards the Nash equilibrium as $\tau \rightarrow \infty$ for this game. Admittedly, it may be necessary to increase τ well beyond the ranged tested herein to observe convergence to the equilibrium, but it may also be the case that equilibrium in this game cannot be reached by strategic reasoning alone. For the equilibrium strategies to be realized, it may be necessary for the players to learn. However, given that the D-Day game occurs in a one-shot setting, learning from repeated play cannot be realized. If this

holds true, the assumption of perfect rationality would severely inhibit analysis and lead to incorrect conclusions.

From an operational planning perspective, the CH analysis indicates that a boundedly rational subordinate commander might struggle to gain an advantage over the enemy in this situation. Therefore, if conditions permit the centralization of this targeting task, the results of Section 5.3 illustrate how knowledge regarding the enemy's strategic thinking ability can be leveraged to develop a strategy that balances planning between enemy capability and intent.

First Strike Decision: Preemptive War Game.

Schelling (1960) contrasted Cold War interaction between the United States and the Soviet Union with regard to preemptive strikes via the “reciprocal fear of surprise attack’. The Prisoner’s Dilemma and Assurance games have been described as two special cases of this dynamic (Kydd, 2007). However, such Phase 0 operations in the Continuum of Military Operations are not a relic of the Cold War. They have continued into modernity (e.g., the Korean Demilitarized Zone and worldwide counterterror operations) and, in this example, we consider a higher headquarters analyzing the effect of a course of action giving rise to such a scenario.

More formally, we illustrate one-shot behavior for both the Prisoner Dilemma and Assurance games with the CH model, and we demonstrate how descriptive EWA-parameter profiles can be captured from observed play with the Prisoner’s Dilemma game. The example payoff structures of both games utilized in this analysis are presented in Tables 21 and 22.

The Prisoner’s Dilemma variant has a single Nash equilibrium of mutual attack that can be found by the iterated removal of dominated strategies. As such, the CH algorithm will converge to this point as $\tau \rightarrow \infty$. This result can be seen in Figure

Table 21. Preemptive War: Prisoner's Dilemma Game

	No Attack	Attack
No Attack	(3,3)	(-1,4)
Attack	(4,-1)	(0,0)

Table 22. Preemptive War: Assurance Game

	No Attack	Attack
No Attack	(8,8)	(1,3)
Attack	(3,1)	(2,2)

29. As in the Nuclear Crisis game, the symmetry of the payoffs enables this graph to describe each player's behavior. However, the pure strategy equilibrium is not reached until τ is approximately five. At lower levels of thinking, the probability of one side choosing not to attack may be non-trivial. Therefore, a unilateral attack is conceivable under bounded rationality, whereas this is not the case under perfect rationality analysis. These results echo those of the brinkmanship example; the probability of a collectively worse outcome (i.e., a bilateral attack) increases with the level of thinking.

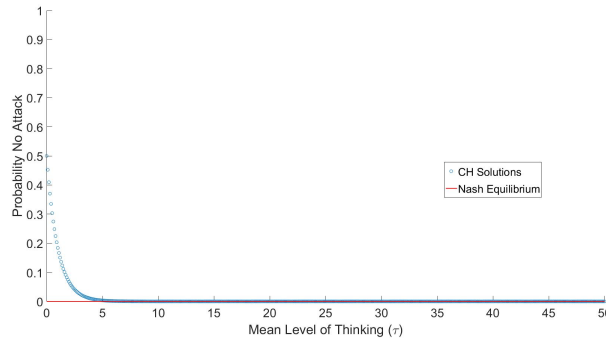


Figure 29. Probability No Attack in Prisoner's Dilemma as a function of τ

Conversely, the Assurance game variant has three Nash equilibrium strategies. Two pure strategy equilibrium with both sides choosing to attack or not attack, and one mixed strategy equilibrium with both players attacking with probability of $\frac{5}{6}$. The

CH algorithm, in this case, shows very encouraging results. In Figure 30, it can be observed that as τ increases, the players both converge (due to symmetry) to a 100% probability of not attacking. Therefore, according to the CH model, the Pareto optimal equilibrium is met as players collectively begin to think more strategically. This is an interesting result since analysis under perfect rationality is unable to differentiate between these equilibrium without additional assumptions (e.g., signaling).

The CH model's results for the Assurance game provide a foil to those of the Prisoner's Dilemma game; the probability of a collectively better outcome increases with the level of thinking. Thus, these results illustrate how BGT methods can begin to quantitatively model individual deterrence situations contingent on the characteristics of the respective leaders. Furthermore, the Assurance game shows convergence towards the referenced equilibrium when τ is approximately five, akin to the aforementioned Prisoner's Dilemma variant. These two variants of the Preemptive War game stand in contrast to the Nuclear Crisis game which continued to demonstrate oscillatory behavior at $\tau = 5$.

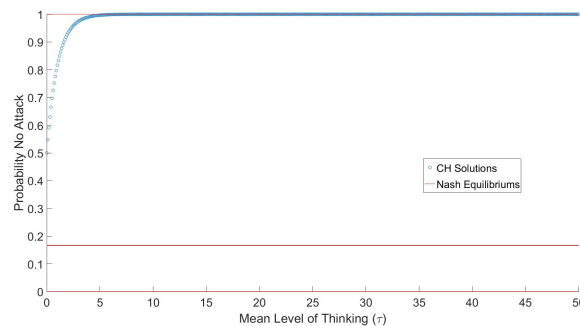


Figure 30. Probability No Attack in Assurance game as a function of τ

We now turn our attention to the descriptive characterization of players with the EWA model from empirical data. In the Nuclear Crisis game, we demonstrated in Section 5.3 how if the appropriate parameters can be estimated for a new game (i.e., players are assumed to play one game analogous to another), future learning behavior can be estimated utilizing the EWA model. However, if instead empirical behavior

has been recorded, the parameters for the game in question can be estimated directly. Although the gathering of such data is a non-trivial endeavor, it may be facilitated if a planner is willing to utilize either noisy data of past conflicts or data that is sufficiently representative of the operational application.

Andreoni and Miller (1993) performed experiments wherein undergraduate students repeatedly played the Prisoner’s Dilemma game. Herein, we assume historical data is assembled such that their *strangers* data set is representative of mean play over ten time periods for the friendly and enemy players in a Prisoner’s Dilemma game as in Table 21. Assuming common parameters across players, we utilize a 50-point spherical packing design on the relevant parameters which we assume are identical for each player: $\phi_i, \delta_i, \kappa_i, \lambda_i$, and $A_i^1(0)$. The respective bounds for these parameters are set at $[0, 1]$, $[0, 1]$, $[0, 1]$, $[0.1, 2.1]$, and $[0.1783, 0.2554]$. Of note, the bounds of $A^1(0)$ are selected to ensure the associate probability at time zero is in the range $[0.3, 0.4]$ when utilizing equation (40). For the sake of illustration, we assume $N(0)$ is known and equals one. Figure 31 presents the results of the best fit parameter setting, namely $\phi_i = 1, \delta_i = 0.4643, \kappa_i = 0.1404$, and $A_i^1(0) = 0.2554$. The quality of fit is via the sum of squared errors on the mean “No Attack” strategy over 200 iterations of the simulation.




Figure 31. Empirical vs. simulated mean for Prisoner Dilemma fitted EWA

As can be observed, the simulation appears to achieve an adequate fit with a total sum of squared error value of 0.0051. However, if the CH model is assumed to

generate initial conditions, an $A_i^1(0)$ value of 0.4 corresponds to a $\tau = 0.22$. This result appears to indicate that, in the experiment conducted by Andreoni and Miller (1993), participants began by selecting strategies nearly randomly. Such results are not necessarily indicative of operational planning problems, but they do demonstrate how insights with regard to levels of strategic thought can be gained by observing empirical data.

Incorporating Behavioral Uncertainty.

Unless a planner is certain of the underlying behavioral parameters which describe a specific interaction, BGT does not provide deterministic results. It shares this characteristic with standard game theory when multiple equilibria exist.

Since equilibriums are merely fixed points, mathematically we cannot determine which equilibrium is most likely to occur. However, this conundrum does not necessarily exist when utilizing BGT solution concepts. That is, if a planner is able to quantify the behavioral uncertainty in a model, quantitative confidence measures of the resulting interactions can be generated. Herein, we discuss the ramifications of such an approach when a higher headquarters is exercising either decentralized or centralized control.

Consider the CH model as utilized in the previous examples wherein the planners contemplate the effect of a particular course of action on the interaction between a subordinate friendly commander and an enemy commander. Therein, we adopted a naive approach with respect to behavioral uncertainty and assumed that $\tau \in \{0, 0.1, \dots, 50\}$ with no additional information. Under these conditions, the brinkmanship example of Section 5.3 bounds the probability either agent deescalates between [0.5,0.9]. If probabilistic information is available on the distribution of τ , this information can be further leveraged to create a probability distribution on the respective players' mixed

strategies, thereby quantifying the behavioral uncertainty. For example, assume τ is distributed consistent with the results of Camerer et al. (2004) with a mean equal to 1.5, and the majority of its mass between [1,2]. Such a distribution could take the form of a simple piecewise uniform distribution as follows):

$$p(\tau) = \begin{cases} \frac{5}{10000}, & \tau \in \{0, 0.01, \dots, 0.99\}, \\ \frac{9}{1010}, & \tau \in \{1, 1.01, \dots, 2\}, \\ \frac{5}{480000}, & \tau \in \{2.01, 2.02, \dots, 50\}. \end{cases} \quad (41)$$

This information can be leveraged to create probability mass functions of mixed strategies for the players (rounded to the second decimal) as in Figure 32. By explicitly recognizing and modeling the stochasticity inherent in the interaction, the resulting explanations derived from the game theoretic analysis mathematically acknowledge that empirical deviations may occur and provide expected magnitudes of these deviations. Moreover, they can be constructed for all relevant actors to provide planners with a more informed estimate of a course of action's effects.

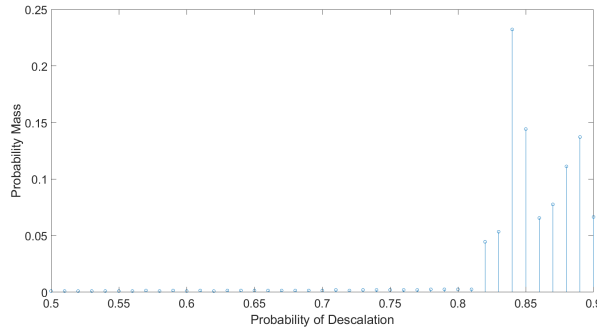


Figure 32. Row Player Probability Mass Function for Likelihood of Deescalation

Furthermore, such characterizations of behavioral uncertainty are extremely useful when the higher headquarters exercises centralized control over a given task. That is, they enable a balance to be obtained between planning against enemy capability

and planning against enemy intent. Haywood (1954) showed that planning against an enemy's capability is an inherently conservative approach akin to a maximin strategy. When utilizing a capability-based approach, "a commander is required to assume that the enemy can discover his decision and will adopt the most effective strategy in opposition". Conversely, planning against an enemy's intent requires a commander to anticipate the adversary's decision and exploit it. If the commander's assessment of the enemy is correct, they may achieve a greater reward than a capability-based approach; however, if the assessment is incorrect, they may also suffer a greater defeat. Therefore, in what follows we describe three mixed approaches to the capability- and intent-based approaches depending upon available information. The uncertainty around the enemy's intent is quantified, and an action is taken to maximize the minimum payoff across this uncertainty.

Consider the D-Day example, but assume the perspective of the Allies. Since the "best response [CH] dictates corresponds to what the highest-step thinkers do", the Allies can identify their preferred strategy by calculating the actions of an arbitrarily large M -step thinker. If the Allies accept that their estimate of τ is subject to error, they may identify robust solutions by quantitatively characterizing this uncertainty. Such characterization may take the form of an uncertainty set (i.e., a range of values), a probability distribution, or an ambiguity set (i.e., a set of probability distributions), and the resulting robust solutions are akin to those derived from the operations research techniques of robust optimization, stochastic programming, and distributionally robust optimization.

Table 23 provides the data for an example of these modeling approaches, assuming τ is in the uncertainty set $\{1, 2, 3, 4, 5\}$. Absent probability information and assuming the payoffs are adequately representative of reality, the Allies should select a robust policy by maximizing their minimum payoff. That is, they would attack

Table 23. Allies' Payoff and Uncertainty Information

τ	Values for Pure Strategies			Probability Distributions		
	$E(\text{Value Calais})$	$E(\text{Normandy})$	$E(\text{Brittany})$	$p_1(\tau)$	$p_2(\tau)$	$p_3(\tau)$
0	0.667	0.417	0.267	0.05	0.10	0.05
1	0.491	0.505	0.355	0.10	0.05	0.05
2	0.271	0.615	0.465	0.50	0.30	0.20
3	0.318	0.499	0.534	0.25	0.25	0.20
4	0.432	0.348	0.570	0.05	0.01	0.40
5	0.629	0.135	0.587	0.05	0.20	0.10

Calais for a maximin payoff of 0.271. Conversely, if the probability distribution $p_1(\tau)$ is accepted as truth, the Allies should attack Normandy for a maximum expected payoff (i.e., expectation over both τ and $E(\pi(s_i^j))$) of 0.528. However, if the Allies believe there is additional error in the probability distribution estimate and consider three distinct variants (i.e., $p_1(\tau)$, $p_2(\tau)$, and $p_3(\tau)$), a distributionally robust policy entails attacking Brittany to attain a maximin expected value of at least 0.472.

Analysis of this form is not unique to the CH model. By characterizing the uncertainty in the EWA model (or any of the models discussed in Section 5.4), similar examinations can be performed. Therefore, BGT does not only provide analysis informed by human psychology, it also allows for the development of more nuanced decision doctrines. Military decisions need not be based on a binary interpretation between capabilities and intent, but can be informed by both.

5.4 Alternative Game Models and BGT Modeling Approaches

The games analyzed in Section 5.3 were considered to ensure their comprehension by a broad readership. However, our choice of illustration does not indicate the inapplicability of BGT to more complicated models. CH can be utilized for games represented in normal form. Scholarship concerning arms races and power transition (e.g. Baliga and Sjöström, 2004; Fearon, 2011, 2018), or those utilizing the Colonel

Blotto framework as a structural basis (e.g. Roberson, 2006; Golman and Page, 2009) are promising venues for future application. Other models of conflict that emphasize bargaining (e.g. Powell, 2015; Slantchev, 2005), especially those assuming imperfect information, may be analyzed with CH when represented in normal form. Such behavioral analysis is especially relevant given the empirical results of LeVeck et al. (2014), who showed the “irrational” behavior of political elites in bargaining situations. The methods described herein are also useful for studying repeated play (e.g. Fearon, 2018), when learning is a foremost consideration. For example, they can be adopted as an alternative to the analysis proposed by Fearon (1994) if time is discretized. Initial game play would be determined utilizing CH, and the subsequent learning process modeled with EWA.

Moreover, we focused primarily on the concepts of step-by-step reasoning and learning in a military operations with normal form (one-shot or repeated) games. However, many other behavioral solution concepts exist for other game structures. For example, the extended form deterrence games of Powell (1990) and Zagare and Kilgour (2000), or the bargaining models of Slantchev (2005) and Powell (2015) can be analyzed utilizing the Agent QRE model of McKelvey and Palfrey (1998). Likewise, alternative EWA models that incorporate dynamic parameters, strategic teaching, or reputation generation can also be considered in repeated play (Camerer et al., 2002). These algorithms can be adapted to include emotional effects, perceptual error, and uncertainty by perturbing the decision rules. For instance, in a manner analogous to trembling hand equilibrium, it may be assumed that the aforementioned factors cause a player to err in their selection of perceived best response strategies. However, to maintain a behaviorist justification, further empirical testing on these adaptations would be required.

Additionally, this research illustrates the need for more empirical work with re-

gard to other considerations in behavioral game theory. For efficient use in military operations planning, it is desirable to empirically test how group deliberation, expertise, and other demographic factors affect the τ -values in the CH model, and/or the parameters in the EWA model. Intuitively, one can postulate that group dynamics affect both the reasoning and learning processes, but there does not currently exist any empirical research to confirm this hypothesis. Moreover, the CH model assumes that τ is population-based but, in a military operations planning setting, the effect of player specific τ -values is a promising area of future work.

Other behavioral constructs informed by cultural understanding are likely imperative in the study of deterrence with regard to the evolved character of 21st Century military operations. Henry Kissinger noted that “the classical notion of deterrence was that there was some consequences before which aggressors and evildoers would recoil. In a world of suicide bombers, that calculation doesn't operate in any comparable way” (Goddard, 2010). However, behavioral game theory may provide the requisite tools to garner insight in this new era of deterrence modeling. For instance, Henrich (200) described how members of the indigenous Machiguenga tribe of Peru behaved decidedly different than players from Los Angeles when playing an Ultimatum game. Camerer (2011) provided an extensive literature review of other findings elucidating this cultural effect, describing how some demographic groups tended to behave in a manner seemingly inconsistent with perfect rationality predictions. However, the author concluded that the behavior originated from a misunderstanding of the player's utility functions and could be described with alternative structures including emotional factors.

Therefore, if a suicide bomber is still governed by a value calculus, there is likely some underlying systematic thought process upon which it is rooted, even if it initially may appear irrational. The underlying logic need not be consistent for it to be

understood, confronted, and deterred (Williams, 2008; Caplan, 2006). This behavioral framework of deterrence focuses on the individual rather than the organizational level, but given the unconventional structure of terrorist groups, this may be a necessary evolution of thought and a productive area of future research.

5.5 Conclusions and Recommendations

This research has taken a first step in integrating behavioral game theory with the planning of military operations. To do so, we reviewed and summarized various bodies of work in BGT, and revisited many classic interstate conflict games to demonstrate how varying insights can be gained through modern, behavioral game theoretic solution concepts. The application of BGT on such fundamental, foundational models depicts in an accessible manner their potential value to military operations planning for an audience having modest or extensive understanding of either military planning or game theory. However, they also indicate the need for future theoretical BGT research, as some more complicated game forms do not yet have behavioral foundations (e.g., stochastic games).

In general, BGT techniques do not supplant the traditional Nash equilibrium solutions but supplement them. As seen in Section 5.3, CH and EWA can provide additional context for games having multiple equilibria by labeling each equilibrium as profiles met through reasoning or learning, respectively. Alternatively, as in the D-Day example, strategic thinking by itself may not allow for an equilibrium profile to be reached, implying that for some interactions occurring only once, insights drawn from perfect rationality analysis are inappropriate and may lead to incorrect conclusions. Finally, as illustrated in Section 5.3, BGT is well-suited to incorporate uncertainty into the underlying analysis, allowing for the development of a balanced capability-and-intent based decision doctrine.

In aggregate, our illustrative instances demonstrate a set of tools that provide utility to inform military operations planning, and provide behavioral game theorists with motivation to study the effects of group dynamics in strategic reasoning and learning for this context. Behavioral game theory removes the deterrent of homo economicus that limited the use of standard game theory in military operations planning, supplanting it with a more representative analytic framework focused on describing actual competitive interaction via the incorporation of psychological factors.

VI. Identifying Behaviorally Robust Maximin Strategies for Normal-form Games under Varying Forms of Uncertainty

Abstract

Recent advances in behavioral game theory address a persistent criticism of traditional solution concepts that rely upon perfect rationality: equilibrium results are often inconsistent with empirical evidence. For normal form games, the Cognitive Hierarchy model is a solution concept based upon a sequential reasoning process, yielding accurate characterizations of experimental human game play. These characterizations are enabled by a statistically estimated parameter describing the average number of reasoning steps players utilize. If an arbitrary player were to know this parameter *ex ante*, they could maximize their expected payoff accordingly. However, given the nature of statistical estimation, such parameter point estimates are unknown prior to experimentation and are susceptible to error afterward. Therefore, we consider the normal form game as a decision problem from the perspective of an arbitrary player who is uncertain of opponents' reasoning ability. Assuming such a player is confronting a set of boundedly rational opponents whose play is characterized by the Cognitive Hierarchy model, we develop a suite of six mathematical programming formulations to maximize the player's minimum payoff, and we identify the appropriate formulation for the level of information regarding an opponent population's reasoning ability. By leveraging robust optimization, stochastic programming, and distributionally robust optimization techniques, our set of models yields prescriptive strategies of play in a normal form game. A software package implementing these constructs is developed and applied to illustrative instances, demonstrating how behaviorally robust strategies vary in accordance with the underlying uncertainty.

6.1 Introduction

Game theory, the study of strategic interaction between self-interested decision-makers, is arguably one of the most significant mathematical constructs of the 20th Century. However, a complete understanding of one of the field's most fundamental questions is paradoxically elusive: In a normal form setting, how should a human agent play a game against human opponents?

An answer to this question depends on the game structure and the agent's beliefs about their opponents. For instance, a Nash equilibrium is a joint strategy profile such that no player can improve their payoff by unilateral deviation. If an agent believes their opponents to be utility maximizers who are mutually consistent (i.e., perfectly rational), then their Nash equilibrium profile can be viewed as a prescriptive strategy of game play (Camerer et al., 2004). Unfortunately, there exists much experimental evidence illustrating that humans are not perfectly rational (e.g., Stahl and Wilson, 1995; Selten, 1998; Costa-Gomes et al., 2001), and that Nash equilibrium profiles are not necessarily predictive of human strategies in games (Camerer, 2011).

Any game theoretic solution technique relying upon perfect rationality may therefore be unable to adequately inform how a human agent should interact with others. While this challenge seems to imply that some other theoretical background (e.g., McKelvey and Palfrey, 1995) must be adapted, the maximin concept is unique among classical game theory in its resilience to the violation of the perfect rationality assumption.

A maximin strategy is one that maximizes an agent's minimum payoff over all possible combinations of opponents' strategies. In many game theoretic texts, the importance of the maximin solution concept is conveyed via its relation to the Nash equilibrium. Namely, by utilizing the results set forth by von Neumann (1928) it can be observed that a maximin strategy in a two-player, zero sum game is a Nash

equilibrium. Similar results have recently been generalized to a special form of n -player games: zero-sum polymatrix games (Cai et al., 2016). However, the maximin concept is interesting in its own right as a robust strategy that guards against a worst case outcome. It does not make any assumption regarding the psychology of an agent's opponents and, for this reason, is resilient to rationality violations. The tradeoff for this resilience is that the maximin strategy may guard against improbable oppositional action and provide an overly conservative strategy.

In operations research, robustness has been utilized as a “decision criteri[on] other than optimality” (Lempert and Collins, 2007). It is useful to do so when model parameters and/or the problem structure is uncertain. As uncertainty regarding either parameters or structure may significantly affect whether an *a priori* identified solution is, in fact, optimal, robustness is used as a safeguard to ensure a solution is prescribed that performs well across a range of problem instantiations.

Of particular interest to this research, solutions resilient to parameter value changes in a mathematical program can be found utilizing robust optimization, stochastic programming, or distributionally robust optimization approaches, depending upon the nature of the parametric uncertainty (Bertsimas et al., 2011). Within a robust optimization (RO) context, the uncertainty associated with parameter values in a mathematical program is characterized via uncertainty sets. These uncertainty sets can assume a variety of forms (e.g., finite, interval-based, polyhedral); however, they all refrain from utilizing probabilistic information (Goerigk and Schöbel, 2016). Generally speaking, research in RO seeks to identify a reformulation of the original mathematical program that incorporates parameteric uncertainty (i.e., a robust counterpart) and provides a “worst-case” solution that is always feasible and performs well with respect to the objective function.

Stochastic programming (SP) is applied in a similar context, but it assumes that

there exists a probability distribution over the uncertainty set. By allowing for infeasibility of some parameteric instantiations, SP methods leverage the probabilistic information via the expected value operator to identify solutions that are of high quality and feasible in expectation.

By contrast, distributionally robust optimization (DRO) can be viewed as a “robust” generalization of SP wherein a set of probability distributions (i.e., an ambiguity set) over the uncertainty set is considered. As with uncertainty sets, ambiguity sets may be characterized in a variety of manners (e.g., moment-based, polytopic) and, akin to RO, DRO guards against a worst-case instantiation. Such a solution is of high quality and feasible in expectation for any possible distribution in the ambiguity set (Gabrel et al., 2014).

In game theory, uncertainty with respect to payoffs has been addressed in multiple ways. For example, games having uncertain payoffs characterized in terms analogous to either RO or SP were explored by Aghassi and Bertsimas (2006) and Harsanyi (1967), respectively. Likewise, uncertainty is considered in games of imperfect information. Solution concepts, in any of these settings, also assume players seek robust strategies.

Conversely, it is assumed herein that the game’s uncertainty derives not from the potential payoffs (i.e., these are assumed known) but from the human opponents’ respective mental frameworks. These adversaries are boundedly rational, and their behavior cannot be predicted perfectly. When playing a normal form game in this setting, it is desirable to adopt a strategy that performs well against such uncertainty. The maximin strategy is one such robust strategy to hedge against opponents’ bounded rationality; however, as previously noted, it may be overly conservative. To address such uncertainty, this work introduces a suite of mathematical programs that collectively modify a player’s maximin strategy to account for variable oppo-

ment rationality by adapting the Cognitive Hierarchy (CH) model of Camerer et al. (2004) to robust optimization, stochastic programming, and distributionally robust optimization settings, alternatively. In this way, we are able to leverage behavioral game theory and multiple optimization sub-disciplines, thereby contributing to the emerging literature of behavioral robustness in games (Brown et al., 2014; Nguyen et al., 2016a,b). The CH model is a behavioral construct describing players by the number of steps of strategic thought they are able to compute of the form described by Schelling (1960): “He thinks we think he thinks we think... he thinks we think he’ll attack; so he thinks we shall; so he will; so we must”. CH is an effective tool for modeling human behavior in normal form games (Camerer et al., 2004; Rogers et al., 2009; Georganas et al., 2015) and is supported by neuroscientific research detailing correlation between brain activity and differing types of strategic thought predicted by the CH model (Bhatt and Camerer, 2005; Camerer et al., 2005; Bhatt et al., 2010; Coricelli and Nagel, 2009).

In this research, we analyze a game not as a holistic system, as is traditionally considered in the literature, but as a decision problem from the perspective of an arbitrary player. This perspective aligns our work with the contemporary focus of the study of games in artificial intelligence, wherein supervised learning techniques are utilized to inform an agent’s preferred strategy (e.g., Moravčík, Matej and Schmid, Martin and Burch, Neil and Lisý, Viliam and Morrill, Dustin and Bard, Nolan and Davis, Trevor and Waugh, Kevin and Johanson, Michael and Bowling, Michael , 2017). However, our research is distinct in that a preferred strategy is developed through behavioral theories versus supervised learning techniques, and uncertainty is explored through uncertainty or ambiguity sets instead of experimentation.

The remainder of this chapter is structured as follows. Section 6.2 reviews the CH modeling framework and formally defines the requisite terminology for the mathemat-

ical programming formulations presented in successive sections. Section 6.3 provides mathematical programming formulations to find maximin strategies in robust optimization, stochastic programming, and distributionally robust optimization settings, respectively, considering either finite or interval-based uncertainty sets for a scenario-defining CH model parameter. Section 6.4 introduces the software package developed for these situations and applies it to multiple games in order to illustrate the effect of uncertainty on behaviorally robust play. Finally, Section 6.5 discusses the implications of this research and its potential for broader application.

6.2 Review of the Cognitive Hierarchy Model

Selten (1998) argued that “the natural way of looking at game situations...is not based on circular concepts, but rather on a step-by-step reasoning procedure”. Furthermore, Schelling (1960) referred to the naturalness of this reasoning structure with regard to the “reciprocal fear of surprise attack”, and he noted that it could be formalized using a potentially infinite set of probabilities composed of the products of the following event probabilities and their reciprocals: P_1 (i.e., the probability the opponent prefers to attack), P_2 (i.e., the probability opponent thinks we prefer to attack), P_3 (i.e., the probability opponent thinks I believe they prefer to attack), and so on. In Schelling’s estimation, “the trouble with the formulation is that nothing generates the series” and that “each probability is an ad hoc estimate” (Schelling, 1960).

Camerer et al. (2004) reinterpreted this concept of step-by-step reasoning into a tractable form with the development of the Cognitive Hierarchy model for normal form games. Their work reinforced the hypothesis of Selten (1998) regarding the human use of sequential logic in strategic interaction via a myriad of experiments illustrating the CH model’s goodness of fit to empirical data from human-subject

testing.

Within the CH modeling framework, players are assumed to be defined both by the number of reasoning steps they use to select an action and by their beliefs regarding the number of steps utilized by other players. It is assumed that players are overconfident; they do not realize their opponents may use an equivalent or greater number of reasoning steps (i.e., a k -step player believes all others are $(k-1)$ -step players or below). The behavior of a group of boundedly rational players is computed recursively starting with 0-step players up to some large k -value (k_{max}). Camerer et al. (2004) assumed 0-step players do not think strategically and adopt a strategy that randomizes uniformly over their action space, 1-step players compute responses assuming all opponents are 0-step thinkers, and so on.

A game is characterized within the CH framework by some true probability distribution on the number of reasoning steps players utilize. The Poisson distribution with mean τ is generally used, and its conventional adoption within CH is supported by the empirical testing conducted by Camerer et al. (2004). A high value of τ indicates players use more reasoning steps, on average. Since players are overconfident, they cannot perceive this true distribution. As such, each k -step player forms their subjective beliefs of encountering an h -step player ($h < k$) as follows:

$$g_k(h) = \frac{f_\tau(h)}{\sum_{l=0}^{k-1} f_\tau(l)} = \frac{\frac{\tau^h e^{-\tau}}{h!}}{\sum_{l=0}^{k-1} \frac{\tau^l e^{-\tau}}{l!}} = \frac{\frac{\tau^h}{h!}}{\sum_{l=0}^{k-1} \frac{\tau^l}{l!}} \quad (42)$$

where $f_\tau(h)$ is the true probability of an h -step player being encountered from the underlying Poisson distribution with rate τ . With this information, a player's action is selected to maximize their expected payoff. More formally, player i chooses an action j such that s_i^j maximizes

$$E_k[\pi(s_i^j)] = \sum_{j'=1}^{m-i} \pi(s_i^j, s_{-i}^{j'}) \left[\sum_{h=0}^{k-1} g_k(h) P_h(s_{-i}^{j'}) \right] \quad (43)$$

where $E_k[\pi(s_i^j)]$ is player i 's expected payoff for playing action j , given player i is a k -step player; $\pi(s_i^j, s_{-i}^{j'})$ is player i 's payoff for selecting action j when his opponents collectively play the action vector j' ; $m-i$ is the number of possible opponent action vectors; and $P_h(s_{-i}^{j'})$ is the probability that h -step players commit to the action vector j' . A k -step player adopts the pure strategy that maximizes their expected payoff. If multiple actions are found to achieve the same maximum payoff, a player is assumed to utilize a strategy that randomizes equally among them.

Therefore, the CH model formalizes the sequential reasoning process into a tractable form by limiting its depth and forcing convergence as k increases. When k grows sufficiently large relative to τ , the actions of a k -step and a $(k+1)$ -step player are indistinguishable; the players have the same beliefs and adopt the same strategy. As such, given a specific τ , the model can be solved by recursively solving equation (43), selecting a best response, and iterating across players until a sufficiently large k -value (k_{max}) is reached. This recursive nature makes the assumed actions of 0-step players foundational to the analysis and results of the CH model. Alternative assumptions for 0-step player actions have been utilized in recent work (e.g., Chong et al., 2016); however, in this research we utilize the assumptions associated with the original CH model and set aside consideration of such alternative frameworks for future research.

6.3 Behaviorally Robust Strategies in Normal Form Games

We utilize the CH model to modify the maximin strategy such that a player of interest can select a high quality strategy given their beliefs regarding the strategic thinking ability of the population from which their opponent(s) derive. As noted by

Camerer et al. (2004), “if the Poisson-CH model is correct for a given game, then the best response the theory dictates corresponds to what the highest-step thinkers do”. Therefore, if τ were known *a priori*, the player of interest would need only to run the CH model and adopt the strategy chosen by an M -step thinker, where M is a sufficiently large integer to observe convergence in successive step thinkers’ strategies.

In practice, τ is likely unknown to a player beforehand since it must be estimated from observed data. Fortunately, Camerer et al. (2004) showed τ to exhibit reasonable regularity patterns such that one can reasonably bound its value *a priori* within some uncertainty set, $U(\tau)$. With this information, an M -step player i no longer associates a scalar payoff $E_M[\pi(s_i^j)]$ with action s_i^j but a set of payoffs for each $\tau \in U(\tau)$, where

$$E_M[\pi(s_i^j), \tau] = \sum_{j'=1}^{m-i} \pi(s_i^j, s_{-i}^{j'}) \left[\sum_{h=0}^{M-1} g_M(h, \tau) P_h(s_{-i}^{j'}, \tau) \right], \quad \forall \tau \in U(\tau), \quad (44a)$$

$$= \sum_{j'=1}^{m-i} \pi(s_i^j, s_{-i}^{j'}) \left[\sum_{h=0}^{M-1} \frac{\tau^h}{h!} P_h(s_{-i}^{j'}, \tau) \right], \quad \forall \tau \in U(\tau), \quad (44b)$$

since player i ’s subjective beliefs about both encountering h -step players, $g_M(h, \tau)$, and their opponents use of the action vector j' , $P_h(s_{-i}^{j'}, \tau)$, depend upon τ .

Consider the column player’s choice in game 10 from the Stahl and Wilson (1995) data set as depicted in Table 24. If it is assumed that this player is unaware of previous testing on the game, and they utilize the point estimate for new games provided by Camerer et al. (2004) of $\tau = 1.5$, their calculated best response is s_{col}^B (i.e., B) because it yields the largest expected payoff of $E_M[\pi(s_{col}^B)] = 35.5$. Conversely, Figure 33 illustrates the expected payoff associated with each action and τ -value should the player assume τ is uncertain but exists within $U(\tau) = \{0, 0.025, 0.05, \dots, 12\}$. Whereas the column player could select an alternative uncertainty set, a subset of the interval $[0, 12]$ is selected for illustration because it contains the τ estimates from Camerer

et al. (2004) across the 12 games tested by Stahl and Wilson (1995).

Table 24. Game 10 from Stahl and Wilson (1995)

	<i>T</i>	<i>M</i>	<i>B</i>
<i>T</i>	(45,45)	(50,41)	(21,40)
<i>M</i>	(41,50)	(0,0)	(40,100)
<i>B</i>	(40,21)	(100,40)	(0,0)

It can be observed from Figure 33 that the consideration of the underlying uncertainty regarding τ complicates the column player's decision. There does not exist a dominating single action across all levels of $\tau \in U(\tau)$. For the set of intervals $\{[0,0.30], [0.625,2.425], [2.925,2.975], [3.475,4.275], [6.1,6.225]\}$ action s_{col}^B is the best response, whereas s_{col}^T yields the greatest expected payoff otherwise.

According to the analysis of Camerer et al. (2004), the empirical data for this game collected by Stahl and Wilson (1995) yields an estimate of $\tau = 11.33$. As shown in Figure 33, a column player selecting a best response of s_{col}^B for a point estimate of $\hat{\tau} = 1.5$ would attain the least expected payoff when $\tau = 11.33$. Therefore, to prescribe a strategy when τ is uncertain, it is desirable to utilize all available information regarding τ to reduce the likelihood of such unfavorable outcomes.

To do so, we first note that, for a known τ -value, player i 's optimal strategy in accordance with the CH model can be written as Problem **Q**:

$$\mathbf{Q} : \max_{P_M(s_i^j)} \sum_{j=1}^{m_i} P_M(s_i^j) E_M[\pi(s_i^j)] \quad (45a)$$

$$\text{subject to} \quad \sum_{j=1}^{m_i} P_M(s_i^j) = 1, \quad (45b)$$

$$P_M(s_i^j) \geq 0, \quad j = 1, \dots, m_i, \quad (45c)$$

wherein the decision variables $P_M(s_i^j)$ determine player i 's strategy, possibly mixed,




Figure 33. *M*-step Column Player Expected Payoffs in Stahl & Wilson Game 10

among the m_i available actions.

With an unknown τ -value, the uncertainty in problem Q stems from two collections of *intermediate* parameters in equation (44a). Both player i 's perceived probability of encountering h -step players, $g_M(h, \tau)$, and the probability player i 's opponents will collectively select action vector j' when utilizing h -steps of thought, $P_h(s_{-i}^{j'}, \tau)$, are uncertain because they are functions of τ . Whereas these terms can be combined into a single uncertain parameter, as depicted in equation (44b), it is useful to represent each intermediate parameter individually for the development of the ensuing mathematical programs since they behave differently across τ -values. More specifically, as discussed in Section 6.3, $g_M(h, \tau)$ is continuous whereas $P_h(s_{-i}^{j'}, \tau)$ is a step function.

In the remainder of this section we illustrate that, if player i can characterize the uncertainty around τ , then behaviorally robust strategies for the game can be developed. Associated with a Poisson distribution, the non-negative parameter τ is

assumed to be bounded above by T , a sufficiently large real number¹, in accordance with the empirical evidence published by Camerer et al. (2004). Moreover, because $\tau \in [0, T]$ is a unidimensional parameter, its associated uncertainty set, $U(\tau)$, is considered to be either finite or interval-based.

For both of these scenarios, we illustrate how varying levels of knowledge about these uncertainty sets can be incorporated into strategy development. We provide RO formulations to select maximin strategies when probabilistic information about τ is unknown, SP formulations when a probability distribution over the uncertainty set is available, and DRO formulations when an ambiguity set of probability distributions is available.

Finite Uncertainty Set over τ .

If $U(\tau)$ is assumed to be some finite subset of $[0, T]$, then Problem Q can be reformulated for RO, SP, and DRO. For such a finite uncertainty, the CH model should be run iteratively for each $\tau \in U(\tau)$ to identify the associated parameters (i.e., $g_k(h, \tau)$ and $P_h(s_{-i}^j, \tau)$) and inform the corresponding optimization model. As such, for a robust optimization framework and a finite uncertainty set $U(\tau)$, Problem Q can be transformed to its robust counterpart Problem **R1**:

$$\mathbf{R1} : \max_{v, P_M(s_i^j)} \quad v \quad (46a)$$

$$\text{subject to} \quad \sum_{j=1}^{m_i} P_M(s_i^j) E_M[\pi(s_i^j), \tau] \geq v, \quad \forall \tau \in U(\tau), \quad (46b)$$

$$\sum_{j=1}^{m_i} P_M(s_i^j) = 1, \quad (46c)$$

$$P_M(s_i^j) \geq 0, \quad j = 1, \dots, m_i, \quad (46d)$$

¹ M should be selected such that $E_k[\pi(s_i^j), \tau]$ has approximately converged in k when $\tau = T$ (i.e., $f_\tau(k)$ are sufficiently small for $k > M$).

wherein v represents the minimum payoff associated with a strategy over all $\tau \in U(\tau)$. The formulation is linear and, when solved to optimality, identifies a strategy $P_M(s_i^j)$ that maximizes the lower bound v .

However, Problem R1 does not incorporate any probabilistic information regarding $U(\tau)$ and, in situations where this information is available, Problem R1 may identify solutions which are overly conservative. For example, Camerer et al. (2004) found that $\tau \in [1, 2]$ characterized empirical behavior in a large density of their experiments. If this information is utilized by player i to inform the estimation of a probability distribution F over $U(\tau)$, then a behaviorally optimized strategy can be found via SP by solving Problem **S1** where \mathbb{E}_F is the expectation over $\tau \sim F$.

$$\mathbf{S1} : \max_{P_M(s_i^j)} \sum_{j=1}^{m_i} P_M(s_i^j) \mathbb{E}_F [E_M[\pi(s_i^j), \tau]] \quad (47a)$$

$$\text{subject to} \quad \sum_{j=1}^{m_i} P_M(s_i^j) = 1, \quad (47b)$$

$$P_M(s_i^j) \geq 0, \quad j = 1, \dots, m_i. \quad (47c)$$

This formulation retains the general structure of R1 but leverages the probabilistic information regarding τ with the expectation operator \mathbb{E}_F . For notational consistency, the domain of $P_M(s_i^j)$ remains unchanged; however, we note that an optimal solution will exist in pure strategies, i.e., where $P_M(s_i^{j^*}) = 1$ for $j^* \in \operatorname{argmin}_{j=1, \dots, m_i} \{E_M[\pi(s_i^j), \tau]\}$ and 0 otherwise.

If player i has a high degree of confidence in the probability distribution F , then the strategy attained by optimally solving Problem S1 should be utilized. Unfortunately, the formulation of Problem S1 requires a single estimate F of the true distribution of τ over $U(\tau)$, and it does not address the effects of estimation error.

To safeguard against estimation error, player i can instead consider an ambiguity set \mathcal{F} of distributions. Empirical results from the work of Camerer et al. (2004) can be leveraged to create a moment-based ambiguity set such as

$$\mathcal{F} := \left\{ F : \sum_{\tau \in U(\tau)} \tau p_{\tau} = c_1, \sum_{\substack{\tau \in U(\tau) \\ c_3 \leq \tau \leq c_4}} p_{\tau} \geq c_2 \right\}, \quad (48)$$

wherein p_{τ} is the probability that some τ in $U(\tau)$ is the true Poisson distribution parameter, c_1 is the identified mean of the unknown distribution (e.g., 1.5 as referenced previously), and $c_2 \in [0, 1]$ is some pre-specified, lower cumulative density of τ between $[c_3, c_4]$ (e.g., $[1, 2]$ as previously discussed). Equipped with such an ambiguity set, a DRO variant of Problem Q having a finite set $U(\tau)$ can be constructed by finding the distributionally robust counterpart of the bilevel program **DR** of the form

$$\begin{aligned} \mathbf{DR} : \max_{P_M(s_i^j)} \min_{F \in \mathcal{F}} & \quad \mathbb{E}_F \left[\sum_{j=1}^{m_i} P_M(s_i^j) E_M[\pi(s_i^j), \tau] \right] \\ & \quad \sum_{j=1}^{m_i} P_M(s_i^j) = 1, \\ & \quad P_M(s_i^j) \geq 0, \quad j = 1, \dots, m_i. \end{aligned} \quad (49a)$$

Solving DR optimally yields the best strategy for player i , assuming the worst possible distribution $F \in \mathcal{F}$ for each feasible strategy. A moment based ambiguity set, as described in equation (48), is a type of polyhedral uncertainty as defined by Bertsimas et al. (2011) and, since both the objective function and the ambiguity set constraints are linear, the bilevel formulation can be reformulated into a single-level

formulation by taking the dual of the lower-level minimization problem. Because $U(\tau)$ is a finite set, the lower-level problem can be alternatively represented as

$$\begin{aligned}
& \min_{p_\tau} && \sum_{\tau \in U(\tau)} p_\tau \left[\sum_{j=1}^{m_i} P_M(s_i^j) E_M[\pi(s_i^j), \tau] \right] \\
& \text{subject to} && \sum_{\tau \in U(\tau)} \tau p_\tau = c_1, \\
& && \sum_{\substack{\tau \in U(\tau) \\ 1 \leq \tau \leq 2}} p_\tau \geq c_2, \\
& && \sum_{\tau \in U(\tau)} p_\tau = 1, \\
& && p_\tau \geq 0, \quad \forall \tau \in U(\tau),
\end{aligned}$$

and its dual substituted into the original program to obtain the DRO variant of Problem Q for a finite uncertainty set $U(\tau)$, formulated as Problem **DR1**.

$$\begin{aligned}
\mathbf{DR1} : & \max_{P_M(s_i^j), z_r} && c_1 z_1 + c_2 z_2 + z_3 \\
& \text{subject to} && \tau z_1 + z_2 + z_3 \leq \sum_{j=1}^{m_i} P_M(s_i^j) E_M[\pi(s_i^j), \tau], \quad \tau \in [c_3, c_4], \\
& && \tau z_1 + z_3 \leq \sum_{j=1}^{m_i} P_M(s_i^j) E_M[\pi(s_i^j), \tau], \quad \tau \notin [c_3, c_4], \\
& && \sum_{j=1}^{m_i} P_M(s_i^j) = 1, \\
& && P_M(s_i^j) \geq 0, \quad j = 1, \dots, m_i, \\
& && z_2 \geq 0, \\
& && z_1, z_3, \text{ unrestricted.}
\end{aligned}$$

This DRO formulation complements the RO and SP variants, allowing a player i to form a behaviorally robust strategy, depending on the level of uncertainty with respect to the parameter τ . The formulations are linear and can be solved to optimality using any number of available commercial optimization solvers. However, these features are made possible by the assumption that $U(\tau)$ is a finite, countable set.

Should the assumption of a finite uncertainty set over $\tau \in [0, T]$ not hold and the uncertainty set is instead interval-based, two options exist to identify a CH-based strategy for player i . First, the methods for a finite uncertainty set can be leveraged as heuristics. Subject to the availability of computational resources, a high granularity of finite values for $\tau \in [0, T]$ can be examined within Problems R1, S1, or DR1, accordingly. We implement such a procedure in the BRMaximin toolbox discussed in Section 6.4. Alternatively, Section 6.3 presents RO, SP, and DRO formulations to directly accommodate an interval-based uncertainty set, $U(\tau)$.

Interval Uncertainty Set over τ .

If $U(\tau)$ is an interval-based uncertainty set, analogous mathematical programming representations can be created to form behaviorally robust strategies. However, the finer fidelity of $U(\tau)$ is accompanied by an increase in complexity both to formulate and solve the corresponding math programs.

We first examine how to create a behaviorally robust strategy within a RO framework by considering the objective function of Problem Q,

$$\sum_{j=1}^{m_i} P_M(s_i^j) E_M[\pi(s_i^j), \tau] = \sum_{j=1}^{m_i} P_M(s_i^j) \sum_{j'=1}^{m_{-i}} \pi(s_i^j, s_{-i}^{j'}) \left[\sum_{h=0}^{k-1} g_k(h, \tau) P_h(s_{-i}^{j'}, \tau) \right].$$

Of the four components on the right-hand side, $P_M(s_i^j)$ is a decision variable, $\pi(s_i^j, s_{-i}^{j'})$

is a constant, and both $g_k(h, \tau)$ and $P_h(s_{-i}^{j'}, \tau)$ are uncertain intermediate parameters of the primary uncertainty τ . As such, a solution to the RO variant of Problem Q for an interval-based uncertainty set $U(\tau)$, is determined by the product of the functions $g_k(h, \tau)$ and $P_h(s_{-i}^{j'}, \tau)$.

Based on the assumptions of Camerer et al. (2004), for given values of k and h , $g_k(h, \tau)$ is a rational function (i.e., the quotient of two polynomials) on $\tau \in U(\tau)$ and, since the denominator is always defined for $\tau > -1$, is continuously differentiable on $U(\tau)$, though it need not be convex. Moreover, the probability player i assigns to each opponent action vector j' is

$$P_h(s_{-i}^{j'}, \tau) = \prod_{d \neq i} P_h(s_d^{j'_d}, \tau),$$

where j'_d is the action of each opponent d in the vector j' , and $P_h(s_d^{j'_d}, \tau)$ is the probability any player d utilizing h -levels of thought ($h < M$) assigns to action j'_d . Per the assumptions of Camerer et al. (2004), $P_h(s_d^{j'_d}, \tau)$ can only take on a finite number of values; more formally, $P_h(s_d^{j'_d}, \tau) : U(\tau) \mapsto \{0, \frac{1}{m_d}, \frac{1}{m_d-1}, \dots, 1\}$. This implies that $P_h(s_d^{j'_d}, \tau)$ is a simple function. Likewise, since the product of simple functions is also a simple function, the same results hold for $P_h(s_{-i}^{j'}, \tau)$.

Therefore, $E_M[\pi(s_i^j), \tau]$ is not guaranteed to be smooth or convex on $U(\tau)$ because it is composed of the sumproduct of $g_k(h, \tau)$ and $P_h(s_{-i}^{j'}, \tau)$ which are rational and simple functions on $U(\tau)$, respectively. If only pure strategies were to be considered, algorithms for non-smooth, non-convex unconstrained optimization, such as gradient sampling (e.g., Burke et al., 2018) or quasi-Newton methods (e.g., Lewis and Overton, 2013), could be applied to each action, $j = 1, \dots, m_i$, yielding the minimum values for each $E_M[\pi(s_i^j), \tau]$ over $U(\tau)$. A RO solution in this case would be the action j^* with the greatest minimum value of $E_M[\pi(s_i^j), \tau]$.

Although such a procedure yields a behaviorally robust strategy, its solution is a

special case of the RO variant of Problem Q having interval-based uncertainty since only pure strategies are considered feasible.

A general RO formulation, that allows for mixed strategies can be modeled incumbent upon $P_h(s_{-i}^{j'}, \tau)$ being not just a simple function but also a step function having values $P_h^\mu(s_{-i}^{j'})$ over a finite number γ of intervals that cover $U(\tau)$, such that $U_\mu(\tau)$ is the closure of any such interval μ . By considering the closure we relax the assumption that opponents mix equally over actions that yield an equivalent maximum payoff, and consider alternative supports wherein a probability of zero may be given to some combinations of these actions. Therefore, the resulting formulation is less sensitive to the CH assumptions and is more tractable computationally.

In this way, the RO variant of Problem Q with an interval-based $U(\tau)$ can be written as

$$\begin{aligned}
& \max_{v, P_M(s_i^j)} && v \\
\text{subject to} &&& \sum_{j=1}^{m_i} P_M(s_i^j) = 1, \\
&&& \left[\min_{\tau_\mu \in U_\mu(\tau)} \sum_{j=1}^{m_i} P_M(s_i^j) \sum_{j'=1}^{m-i} \sum_{h=0}^{M-1} \pi(s_i^j, s_{-i}^{j'}) g_M(h, \tau) P_h^\mu(s_{-i}^{j'}) \right] \geq v, \\
&&& \mu = 1, \dots, \gamma, \tag{50a} \\
&&& P_M(s_i^j) \geq 0, \quad j = 1, \dots, m_i.
\end{aligned}$$

However, since $g_M(h, \tau)$ is continuous over $U(\tau)$, it is also continuous over each $U_\mu(\tau)$. Call $\Delta_1^{h\mu}$ and $\Delta_2^{h\mu}$ the minimum and maximum values of $g_M(h, \tau)$ in some interval $U_\mu(\tau)$ for some h . These can be found using any number of unconstrained optimization techniques (e.g., line search) for continuous functions. By the Intermediate Value Theorem, the resulting interval-based uncertainty set of $g_M(h, \tau)$ for each

instantiation of constraint (50a) is $[\Delta_1^{h\mu}, \Delta_2^{h\mu}]$. Therefore, we can consider $g_M(h, \tau)$ directly by introducing the decision variables $g_M^\mu(h)$ to represent their worst-case values in each interval $U_\mu(\tau)$. Each optimization problem embedded within constraint (50a) can be alternatively represented as

$$\min_{g_M^\mu(h)} \sum_{j=1}^{m_i} P_M(s_i^j) \sum_{j'=1}^{m-i} \sum_{h=0}^{M-1} \pi(s_i^j, s_{-i}^{j'}) g_M^\mu(h) P_h^\mu(s_{-i}^{j'})$$

$$\text{subject to} \quad g_M^\mu(h) \geq \Delta_1^{h\mu}, \quad h = 0, \dots, M-1, \quad (51a)$$

$$-g_M^\mu(h) \geq -\Delta_2^{h\mu}, \quad h = 0, \dots, M-1, \quad (51b)$$

and, by taking the dual of each such formulation, in a manner similar to Problem DR1, we can express the robust counterpart of Problem Q with an interval-based uncertainty set $U(\tau)$ as

$$\begin{aligned} \mathbf{R2} : \quad & \max_{v, P_M(s_i^j), z_1^{h\mu}, z_2^{h\mu}} v \\ & \text{subject to} \quad \sum_{j=1}^{m_i} P_M(s_i^j) = 1, \\ & \sum_{h=0}^{M-1} \Delta_1^{h\mu} z_1^{h\mu} - \Delta_2^{h\mu} z_2^{h\mu} \geq v, \quad \mu = 1, \dots, \gamma, \end{aligned} \quad (52a)$$

$$\begin{aligned} z_1^{h\mu} - z_2^{h\mu} &= \sum_{j=1}^{m_i} \sum_{j'=1}^{m-i} P_M(s_i^j) \pi(s_i^j, s_{-i}^{j'}) P_h^\mu(s_{-i}^{j'}), \\ & \mu = 1, \dots, \gamma, h = 0, \dots, M-1, \end{aligned} \quad (52b)$$

$$z_1^{h\mu}, z_2^{h\mu} \geq 0, \quad \mu = 1, \dots, \gamma, h = 0, \dots, M-1, \quad (52c)$$

$$P_M(s_i^j) \geq 0, \quad j = 1, \dots, m_i,$$

wherein $z_1^{h\mu}$ and $z_2^{h\mu}$ are the dual variables associated with constraints (51a) and

(51b), respectively. This single level reformulation is linear and can be readily solved utilizing a commercial solver.

However, as previously mentioned, the form of Problem R2 depends upon $P_h(s_{-i}^{j'}, \tau)$ being a step function, a condition proven to hold in Theorem 6.3.1.

For notational ease, we define $\Phi_{hi}(\tau)$ to be the upper envelope of the collection of functions $\{E_h[s_i^1, \tau], \dots, E_h[s_i^{m_i}, \tau]\}$ for a given player i and level of thought $h < M$ as a function of τ , or

$$\Phi_{hi}(\tau) = \max_j E_h[\pi(s_i^j, \tau)] . \quad (53)$$

As such, $P_h(s_i^j, \tau) > 0$ for a given $\bar{\tau} \in U(\tau)$ if and only if $j \in \operatorname{argmax}\{E_h[\pi(s_i^j), \bar{\tau}]\}$, i.e., due to the assumptions of Camerer et al. (2004) regarding the selection of perceived best responses.

Theorem 6.3.1 *For any given action vector j' and $h < M$, $P_h(s_{-i}^{j'}, \tau)$ is a step function in τ .*

Proof For any player d , action w , and $h = 0$, $P_0(s_d^w, \tau) = \frac{1}{m_d} \forall \tau \in U(\tau)$, and is trivially a step function by the assumptions of Camerer et al. (2004). Therefore, so is $P_0(s_{-i}^{j'}, \tau) = \prod_{d \neq i} P_0(s_d^{j'_d}, \tau)$.

Moreover, for any player d and $h = 1$

$$\Phi_{1d}(\tau) = \max_w E_1[\pi(s_d^w, \tau)] = \max_w \left[\sum_{w'=1}^{m-d} \pi(s_d^w, s_{-d}^{w'}) \right] ,$$

implying that $\Phi_{1d}(\tau)$ has no break points, $P_h(s_d^w, \tau)$ is a step function for $w = 1, \dots, m_d$, and $P_1(s_{-i}^{j'}, \tau)$ is a step function as well.

Assume that $P_h(s_d^w, \tau)$ is step function for $h = 2, \dots, n$. Then so is $P_h(s_{-d}^{w'}, \tau) = \prod_{\nu \neq d} P_h(s_{\nu}^{w'_\nu}, \tau)$, and it has corresponding set of intervals of $U_y(\tau)$, $y = 1, \dots, \lambda_y$ covering $U(\tau)$. Given $g_k(h, \tau)$ is a continuous, rational function on $U(\tau)$ for a given h , at $h = n + 1$ we have

$$E_{n+1}[\pi(s_d^w), \tau] = \sum_{w'=1}^{m-d} \sum_{h=0}^n \pi(s_d^w, s_{-d}^{w'}) g_{n+1}(h, \tau) P_h(s_{-d}^{w'}, \tau)$$

piecewise continuous in $U(\tau)$, and a rational function on each $U_y(\tau)$.

Call \mathcal{D} the set of degenerate $U_y(\tau)$ intervals. If $\tau \in U_y(\tau) \subseteq \mathcal{D}$, then $\Phi_{hd}(\tau)$ is found via equation (53) and, as such, $P_{n+1}(s_d^w, \tau) = \beta_d^w(y) = \frac{1}{|W^*|}$ for $w \in W^*$ and zero otherwise, where W^* is the set of actions with an expected payoff defining $\Phi_{hd}(\tau)$.

Call \mathcal{N} the set of non-degenerate $U_y(\tau)$ intervals. If $\tau \in U_y(\tau) \subseteq \mathcal{N}$, we note the following. For any two actions, w_1 and w_2 , we can write $E_{n+1}[\pi(s_d^{w_1}), \tau] = \frac{D_1(\tau)}{D_2(\tau)}$ and $E_{n+1}[\pi(s_d^{w_2}), \tau] = \frac{D_3(\tau)}{D_4(\tau)}$, where $D_\delta(\tau)$ are each polynomials of order \dot{D}_δ . The points of intersection of these two functions can be found by solving

$$E_{n+1}[\pi(s_d^{w_1}), \tau] - E_{n+1}[\pi(s_d^{w_2}), \tau] = \frac{D_1(\tau)D_4(\tau) - D_2(\tau)D_3(\tau)}{D_2(\tau)D_4(\tau)} = 0.$$

$D_2(\tau)D_4(\tau)$ has no zeros in $U(\tau)$ since $g_k(h)$ is always defined in this domain. $D_1(\tau)D_4(\tau) - D_2(\tau)D_3(\tau)$ is another polynomial with order no greater than $\max\{\dot{D}_1 + \dot{D}_4, \dot{D}_2 + \dot{D}_3\}$ by the Fundamental Theorem of Algebra. Assuming $E_{n+1}[\pi(s_d^{w_1}), \tau] - E_{n+1}[\pi(s_d^{w_2}), \tau]$ does not equal the zero function (i.e., $w_1 \neq w_2$), then it has a maximum number of zeros equal to the order of $D_1(\tau)D_4(\tau) - D_2(\tau)D_3(\tau)$.

Therefore, within any $U_y(\tau) \in \mathcal{N}$, there exists a finite number of intersections among the set of functions $\{E_{n+1}[\pi(s_d^w), \tau]\}$, implying $\Phi_{(n+1)d}(\tau)$ has a finite number of intervals within which some set of actions W^* defines the upper envelope and $P_{n+1}(s_d^w, \tau)$ is some constant. That is, $\forall U_y(\tau) \in \mathcal{N}$, there exists a finite number λ_y of intervals $U_{yr}(\tau)$ that cover $U_y(\tau)$ and wherein each $P_{n+1}(s_d^w, \tau)$ equals some constant $\beta_d^w(y, r)$

Collectively, these results allow $P_{n+1}(s_d^w, \tau)$ to be written as

$$P_{n+1}(s_d^w, \tau) = \left[\sum_{U_y(\tau) \in \mathcal{N}} \sum_{r=1}^{\lambda_y} \beta_d^w(y, r) \mathbb{I}_{yr} \right] + \left[\sum_{U_y(\tau) \in \mathcal{D}} \beta_d^w(y) \mathbb{I}_y \right]$$

wherein \mathbb{I}_y and \mathbb{I}_{yr} are indicator functions for the intervals $U_y(\tau)$ and $U_{yr}(\tau)$, respectively. Therefore, $P_{n+1}(s_d^w, \tau)$ is a step function, and so is their product $P_{n+1}(s_{-i}^{j'}, \tau) = \prod_{d \neq i} P_{n+1}(s_d^{j'}, \tau)$. \square

In this way, the uncertainty pertaining to intermediate parameter $P_h(s_{-i}^{j'}, \tau)$ can be handled explicitly by partitioning $U(\tau)$ into the appropriate intervals wherein it is constant. In practice, these intervals can be found by repeatedly finding the upper envelopes in τ of each player's payoff functions for a given action, opponent action and level of thought, and recording for which τ -values we have $P_h(s_{-i}^{j'}, \tau) > 0$.

The results from Theorem 6.3.1 are also useful in formulating the SP variant of Problem Q with an interval-based $U(\tau)$, which we denote as **S2**:

$$\begin{aligned} \mathbf{S2} : \max_{P_M(s_i^j)} & \sum_{j=1}^{m_i} P_M(s_i^j) \mathbb{E}_F \left[E_M[\pi(s_i^j), \tau] \right] \\ \text{subject to} & \sum_{j=1}^{m_i} P_M(s_i^j) = 1, \\ & P_M(s_i^j) \geq 0, \quad j = 1, \dots, m_i. \end{aligned}$$

We note that S2 has the same general form as S1. However, we can leverage the fact that τ is a continuous random variable and the piecewise form of $P_h(s_{-i}^{j'}, \tau)$ to simplify the objective function per the following theorem.

Theorem 6.3.2 *For τ distributed by an arbitrary, continuous probability distribution function $f(\tau)$ over $U(\tau)$, we have*

$$\mathbb{E}_F \left[E_M[\pi(s_i^j), \tau] \right] \approx \sum_{j'=1}^{m-i} \sum_{h=0}^{M-1} \sum_{\mu=1}^{\gamma} \pi(s_i^j, s_{-i}^{j'}) P_h^\mu(s_{-i}^{j'}) \int_{U_\mu(\tau)} \hat{f}(h, \tau) f(\tau) d\tau$$

where $\hat{f}(h, \tau)$ is a gamma probability density function with rate parameter $\beta = 1$ and shape parameter $\alpha = h + 1$.

Proof By properties of the expectation operator, we can write

$$\mathbb{E}_F \left[E_M[\pi(s_i^j), \tau] \right] = \sum_{j'=1}^{m-i} \pi(s_i^j, s_{-i}^{j'}) \left[\sum_{h=0}^{k-1} \mathbb{E}_F \left[g_k(h, \tau) P_h(s_{-i}^{j'}, \tau) \right] \right]$$

Since $g_k(h, \tau) P_h(s_{-i}^{j'}, \tau)$ is a piecewise continuous function it can be written as

$$g_k(h, \tau) P_h(s_{-i}^{j'}, \tau) = \sum_{\mu=1}^{\gamma} \mathbb{I}_{\mu} \frac{\frac{\tau^h e^{-\tau}}{h!}}{\left(\sum_{l=0}^{M-1} \frac{\tau^l e^{-\tau}}{l!} \right)} P_h^{\mu}$$

where \mathbb{I}_{μ} is the characteristic function associated with the interval $U_{\mu}(\tau)$. The denominator in the summand is the sum of Poisson probability mass functions and, when $M \rightarrow \infty$, it converges toward 1. Therefore, for a sufficiently large M (relative to τ) we can write

$$g_k(h, \tau) P_h(s_{-i}^{j'}) \approx \sum_{\mu=1}^{\gamma} \mathbb{I}_{\mu} \frac{\tau^h e^{-\tau}}{h!} P_h^{\mu},$$

and

$$\begin{aligned} \mathbb{E}_F \left[g_k(h, \tau) P_h(s_{-i}^{j'}, \tau) \right] &\approx \sum_{\mu=1}^{\gamma} P_h^{\mu} \int_{U_{\mu}(\tau)} \frac{\tau^h e^{-\tau}}{h!} f(\tau) d\tau \\ &\approx \sum_{\mu=1}^{\gamma} P_h^{\mu} \int_{U_{\mu}(\tau)} \hat{f}(h, \tau) f(\tau) d\tau, \end{aligned}$$

where $\hat{f}(h, \tau)$ is a gamma probability density function with rate parameter $\beta = 1$ and shape parameter $\alpha = h + 1$. □

The results of Theorem 6.3.2 facilitate the computation of the expected payoff associated with each action j . Since the integral considers the product of two density functions, $\hat{f}(h, \tau)$ and $f(\tau)$, a modeler is potentially able to simplify it. After these integrals are solved and $\mathbb{E}_F \left[g_k(h, \tau) P_h(s_{-i}^{j'}, \tau) \right]$ is found, S2 becomes a simple linear

program with an optimal solution yielding a behaviorally robust strategy in a SP setting.

The combination of Theorems 6.3.1 and 6.3.2 also proves useful for developing distributionally robust maximin strategies with an interval-based $U(\tau)$. Such a problem utilizing the same ambiguity set \mathcal{F} from Section 6.3 can be represented as DR. However, the results of this section allow for the development of a tractable approximate DRO counterpart that converges toward the exact DRO counterpart as parameters $\theta \rightarrow 0$ and $M \rightarrow \infty$.

Utilizing the properties of expectation and Theorems 6.3.1 and 6.3.2, the objective function (49a) can be written as

$$\begin{aligned} \mathbb{E}_F \left[\sum_{j=1}^{m_i} P_M(s_i^j) E_M[\pi(s_i^j), \tau] \right] \\ = \sum_{j=1}^{m_i} P_M(s_i^j) \sum_{j'=1}^{m_{-i}} \pi(s_i^j, s_{-i}^{j'}) \left[\sum_{h=0}^{k-1} \mathbb{E}_F \left[g_k(h, \tau) P_h(s_{-i}^{j'}, \tau) \right] \right] \\ \approx \sum_{j=1}^{m_i} P_M(s_i^j) \sum_{j'=1}^{m_{-i}} \pi(s_i^j, s_{-i}^{j'}) \left[\sum_{h=0}^{k-1} \sum_{\mu=1}^{\gamma} P_h^\mu \int_{U_\mu(\tau)} \hat{f}(h, \tau) f(\tau) d\tau \right]. \end{aligned}$$

We define $U_b(\tau)$ to be a finite set of $\dot{\gamma}$ closed intervals of the form $[a_{b1}, a_{b2}]$ covering $U(\tau)$ that is at least as fine as the set of $U_\mu(\tau)$ (i.e., $\dot{\gamma} \geq \gamma$) such that $a_{11} = \min[U(\tau)]$, $a_{\dot{\gamma}2} = \max[U(\tau)]$, and $a_{b2} = a_{(b+1)1}$ for $b < \dot{\gamma}$. Moreover, we assume each $U_b(\tau)$ is of equal length θ and all endpoints of $U_\mu(\tau)$ correspond to endpoints among the $\dot{\gamma}$ intervals. This property implies that all non-degenerate $U_\mu(\tau)$ interval lengths are multiples of θ (i.e., $\forall \mu \exists n \in \mathbb{N}$ such that $len(U_\mu(\tau)) = n\theta$ where $len(\cdot)$ is the length function).

Once a θ -value is set and $U_b(\tau)$ determined, we utilize them to approximate the ambiguity set \mathcal{F} . That is, we consider all stepwise probability distribution functions

(i.e., histograms) with intervals $U_b(\tau)$. Therefore, it is clear that as $\theta \rightarrow 0$, our approximate ambiguity set converges to \mathcal{F} . Since within each $U_b(\tau)$ the distribution $f(\tau)$ is a constant y_b , we can write

$$\begin{aligned} \sum_{\mu=1}^{\gamma} P_h^{\mu} \int_{U_{\mu}(\tau)} \hat{f}(h, \tau) f(\tau) d\tau &= \sum_{b=1}^{\dot{\gamma}} P_h^{U_b(\tau)} y_b \int_{U_b(\tau)} \hat{f}(h, \tau) d\tau \\ &= \sum_{b=1}^{\dot{\gamma}} P_h^{U_b(\tau)} y_b \Omega(h, U_b(\tau)) \end{aligned}$$

where $P_h^{U_b(\tau)} = P_h^{\mu}$ for $U_b(\tau) \subseteq U_{\mu}(\tau)$, $\Omega(h, U_b(\tau)) = \hat{F}(h, a_{b2}) - \hat{F}(h, a_{b1})$, and $\hat{F}(h, \tau)$ is the gamma cumulative distribution function with $\alpha = h + 1$ and $\beta = 1$ for τ .

With this knowledge, Problem DR can be approximated as

$$\begin{aligned} \max_{P_M(s_i^j)} \min_{y_b} & \left[\sum_{j=1}^{m_i} P_M(s_i^j) \sum_{j'=1}^{m-i} \pi(s_i^j, s_{-i}^{j'}) \left[\sum_{h=0}^{k-1} \sum_{b=1}^{\dot{\gamma}} P_h^{U_b(\tau)} y_b \Omega(h, U_b(\tau)) \right] \right] \\ & \sum_{j=1}^{m_i} P_M(s_i^j) = 1, \\ & P_M(s_i^j) \geq 0, \quad j = 1, \dots, m_i, \\ & \sum_{b=1}^{\dot{\gamma}} \left(\frac{a_{b1} + a_{b2}}{2} \right) y_b = c_1, \\ & \sum_{c_3 \leq a_{b1} \leq c_4} \theta y_b \geq c_2, \\ & \sum_{b=1}^{\dot{\gamma}} \theta y_b = 1, \\ & y_b \geq 0. \end{aligned} \tag{55a}$$

provided that θ has been selected such that c_3 and c_4 are endpoints to some interval $U_b(\tau)$. This qualification is only necessary to ensure the correct calculation of con-

straint (55a) and only relevant when the ambiguity set considers the dispersion of its distributions.

Since all functions are linear, this bilevel program can be reformulated by taking the dual of the lower-level decisionmaker's problem. The resulting single level formulation is denoted as Problem **DR2**.

$$\begin{aligned}
\mathbf{DR2} : \quad & \max_{P_M(s_i^j), z_k} && c_1 z_1 + c_2 z_2 + z_3 \\
\text{subject to} & && \left(\frac{a_{b1} + a_{b2}}{2} \right) z_1 + \theta(z_2 + z_3) \leq \sum_{j=1}^{m_i} P_M(s_i^j) \chi_{ib}^j, \\
& && \forall b : a_{b1} \in [c_3, c_4], \\
& && \left(\frac{a_{b1} + a_{b2}}{2} \right) z_1 + \theta z_3 \leq \sum_{j=1}^{m_i} P_M(s_i^j) \chi_{ib}^j, \\
& && \forall b : a_{b1} \notin [c_3, c_4], \\
& && z_2 \geq 0 \\
& && z_1, z_3, \text{ unrestricted.}
\end{aligned}$$

where

$$\chi_{ib}^j = \sum_{j'=1}^{m-i} \pi(s_i^j, s_{-i}^{j'}) \left[\sum_{h=0}^{M-1} P_h^{U_b(\tau)} \Omega(h, U_b(\tau)) \right]$$

This formulation is an approximate distributionally robust counterpart for the ambiguity set \mathcal{F} that converges to the true counterpart as $\theta \rightarrow 0$ and $M \rightarrow \infty$. Moreover, it allows player i to form a behaviorally robust strategy in a DRO setting with an interval-based $U(\tau)$ via the solution of a linear program.

6.4 Behaviorally Robust Strategies for the Stahl & Wilson Games

The appropriate behaviorally robust strategy varies based upon the information a player has available regarding the parameter τ . This section examines such strategy alterations for the twelve games provided by Stahl and Wilson (1995) and tests them on their human-subject data.

We developed the *BRMaximin* toolbox and utilized it on these examples. This software package includes a suite of MATLAB functions that implements the baseline Cognitive Hierarchy model, plots all $E_M[\pi(s_i^j), \tau]$ functions (i.e., belonging to each action j) for a game across an uncertainty set, and identifies behaviorally robust strategies for a finite or interval-based $U(\tau)$ in a RO, SP and DRO setting, respectively. For each function, the input required consists of a payoff matrix, uncertainty information, and an M -value. For model variants having interval-based uncertainty sets, the BRMaximin toolbox leverages a discrete mesh approximation (with user-defined spacing) over the region for the purpose of computational tractability and, for SP2, uses a beta distribution due to its bounded continuous nature and flexibility of shape. The software package and related documentation are available on both the MATLAB File Exchange² and GitHub³.

The Twelve Stahl & Wilson Games.

The twelve games utilized by Stahl and Wilson (1995) are symmetric matrix games wherein each player has three available actions (i.e., T , M , and B). Game 10 of this set is presented in Table 24. Stahl and Wilson (1995) conducted human-subject testing on each game wherein 48 individuals adopted the role of the row player. Moreover, Camerer et al. (2004) utilized the same data in their original testing of the Cognitive Hierarchy model.

²URL: <https://www.mathworks.com/matlabcentral/fileexchange/71117-brmaximin>

³URL: <https://github.com/caballerow/BRMaximin>

For varying information sets regarding τ and for each of these twelve games, we identify the column player's strategies prescribed by the appropriate model from Section 6.3. We subsequently evaluate the performance of these six strategies, as well as the strategy that would be prescribed by Camerer et al. (2004) under an assumption that the point estimate of $\tau = 1.5$ is correct, against the responses of the 48 individuals tested by Stahl and Wilson (1995). The specifications of the information conditions considered for each game are presented in Table 25.

Table 25. Seven information conditions for Stahl & Wilson games

Information Condition	Assumptions on true τ -value
PE	$\tau = 1.5$
R1	$\tau \in U_1(\tau) = \{0, 0.5, 1, 1.5, 2, 2.5, 3\}$
S1	Distribution F over $U_1(\tau)$; $(p_0, p_{0.5}, p_1, p_{1.5}, p_2, p_{2.5}, p_3) = (0.02, 0.1, 0.15, 0.46, 0.15, 0.1, 0.02)$
DR1	Collection of discrete distribution \mathcal{F} over $U_1(\tau)$ with $c_1 = 1.5$, $c_2 = 0.5$, $c_3 = 1$, and $c_4 = 2$
R2	$\tau \in U_2(\tau) = [1, 12]$
S2	$\tau \sim \text{Beta}(10, 75)$ over $U_2(\tau)$
DR2	Collection of continuous distributions \mathcal{F} over $U_2(\tau)$ with $c_1 = 1.5$, $c_2 = 0.5$, $c_3 = 1$, and $c_4 = 2$

These information conditions are selected to illustrate the effect an agent's knowledge regarding τ can have on their behaviorally robust strategy. The condition PE utilizes the point estimate provided by Camerer et al. (2004) without quantifying the uncertainty from which it was derived. Condition R1 examines the effect of a small, finite uncertainty set; condition S1 considers a probability distribution over this set with mean 1.5; and condition R1 considers all ambiguity sets over the uncertainty with mean 1.5 and at least 50% of their density lying within $[1, 2]$. Conditions R2, S2, and DR2 similarly examine these different forms of uncertainty but on a larger, interval-based uncertainty set for τ . Setting $M = k_{max} = 40$ and using a mesh spacing of 0.025 for the interval-based instances, the resulting behaviorally robust strategies for the BRMaximin toolbox functions are reported in Table 26.

The degree to which the information condition affects the resulting strategy varies by game. In Game 1, every prescribed strategy corresponds to the (unique) maximin strategy; the solutions do not differ by information condition. However, the solutions

Table 26. Prescribed mixed strategies for all Stahl & Wilson games

Game	PE	R1	S1	DR1	R2	S2	DR2
1	(0, 1, 0)*	(0, 1, 0)*	(0, 1, 0)*	(0, 1, 0)*	(0, 1, 0)*	(0, 1, 0)*	(0, 1, 0)*
2	(0, 1, 0)	(0, 1, 0)	(0, 1, 0)	(0, 1, 0)	(0.91, 0.09, 0)	(0, 1, 0)	(0.37, 0.63, 0)
3	(0, 1, 0)	(0, 1, 0)	(0, 0, 1)	(0, 1, 0)	(0, 1, 0)	(0, 1, 0)	(0, 1, 0)
4	(0, 1, 0)	(0.34, 0.66, 0)	(1, 0, 0)	(0.29, 0.71, 0)	(0.62, 0.38, 0)	(1, 0, 0)	(0.33, 0.67, 0)
5	(0, 0, 1)*	(0.75, 0, 0.25)	(0, 0, 1)*	(0, 0, 1)*	(0, 0, 1)*	(0, 0, 1)*	(0, 0, 1)*
6	(0, 1, 0)*	(0.53, 0.47, 0)	(0, 1, 0)*	(0.39, 0.61, 0)	(0, 1, 0)*	(0, 1, 0)*	(0.32, 0.68, 0)
7	(1, 0, 0)	(1, 0, 0)	(1, 0, 0)	(1, 0, 0)	(0, 0.30, 0.70)*	(1, 0, 0)	(0.65, 0, 0.35)
8	(0, 1, 0)	(0, 1, 0)	(0, 1, 0)	(0, 0.10, 0.90)	(0, 1, 0)	(0, 1, 0)	(0, 0, 1)
9	(0, 0, 1)	(0.54, 0, 0.46)	(0, 0, 1)	(0.50, 0, 0.50)	(0.56, 0, 0.44)	(0, 0, 1)	(0.54, 0, 0.46)
10	(0, 0, 1)	(1, 0, 0)	(0, 0, 1)	(0.97, 0, 0.03)	(1, 0, 0)	(0, 0, 1)	(1, 0, 0)
11	(0, 0, 1)	(0, 0, 1)	(0, 0, 1)	(0, 0, 1)	(0.30, 0.59, 0.11)*	(0, 0, 1)	(0, 0.04, 0.96)
12	(1, 0, 0)	(0.90, 0, 0.10)	(1, 0, 0)	(0.36, 0, 0.64)	(1, 0, 0)	(1, 0, 0)	(0.34, 0, 0.66)

* Strategy coincides with Maximin solution

to other games illustrate variability based on uncertainty conditions (e.g., Game 4). One factor driving this behavior derives from the form of the $E_M[\pi(s_{col}^j), \tau]$ functions in each game. Consider each of these functions for games 1, 4, and 10, as graphed in Figures 34, 35 and 33. In both Figures 1 and 3, the expected payoff functions exhibit multiple discontinuities and, within each figure, the upper envelopes of the the set of functions corresponds to different strategies over multiple, disjoint intervals. In contrast, the upper envelope in Figure 34 is defined by a much simpler piecewise continuous function. The best expected outcome is $E_M[\pi(s_{col}^T), \tau]$ from $\tau = [0, 0.125]$ and $E_M[\pi(s_{col}^M), \tau]$ elsewhere. Thus, the information conditions regarding τ have a notable effect on prescribed strategies for Games 4 and 10 and no effect on prescribed strategies for Game 1.

Table 27 reports the worst, average, and best performance of the respective column player strategies prescribed in Table 26 for all 12 games, considering as opponents the responses of the 48 players tested by Stahl and Wilson (1995). For each game, Table 27 also reports the regressed τ -values from Camerer et al. (2004) (i.e., $\hat{\tau}$) as well as the payoff statistics for both the CH model's best response (BR) strategy with the regressed $\hat{\tau}$ and the standard maximin (MM) strategy. These results highlight the effect a player's understanding of the opponent population can have on their expected




Figure 34. *M*-step Column Player Expected Payoffs in Stahl & Wilson Game 1




Figure 35. *M*-step Column Player Expected Payoffs in Stahl & Wilson Game 4

payoff. As some of the column player's prescribed strategies are probabilistic, the results are reported in terms of their expected payoff against this set of opponents.

Table 27. Descriptive expected payoff statistics versus 48 Stahl & Wilson players

Game	$\hat{\tau}$	Statistic	BR	MM	PE	R1	S1	DR1	R2	S2	DR2
1	2.93	Min	40	40	40	40	40	40	40	40	40
		Avg	44.69	44.69	44.69	44.69	44.69	44.69	44.69	44.69	44.69
		Max	65	65	65	65	65	65	65	65	65
2	0	Min	15	41.54	15	15	15	15	37.66	15	24.19
		Avg	60	62.21	60	60	60	60	62.27	60	60.92
		Max	100	74.61	100	100	100	10	74.53	100	79.79
3	1.40	Min	35	35	35	35	35	35	35	35	35
		Avg	45.31	45.21	45.31	45.31	45.14	45.31	45.31	45.31	45.31
		Max	80	92.38	80	80	100	80	80	80	80
4	2.34	Min	30	37	10	36.64	30	35.70	33.83	30	36.68
		Avg	46.46	39.97	37.19	40.30	46.46	39.84	42.91	46.46	40.27
		Max	100	46.50	45	46.68	100	46.43	65.57	100	46.66
5	2.01	Min	20	20	20	12.44	20	20	20	20	20
		Avg	47.71	47.71	47.71	38.11	47.71	47.71	47.71	47.71	47.71
		Max	50	60	60	87.77	60	60	60	60	60
6	0	Min	25	31	31	30.47	31	30.61	31	31	30.69
		Avg	53.65	44.73	44.73	49.44	44.73	48.16	47.73	47.73	47.55
		Max	100	60	60	76.88	60	69.87	60	60	66.50
7	5.37	Min	30	47.04	30	30	30	30	46.99	30	37.08
		Avg	58.96	49.09	58.96	58.96	58.96	58.96	49	58.96	57.37
		Max	100	52.82	100	100	100	100	52.46	100	91.15
8	0	Min	20	42.86	20	20	20	38	20.01	20	40
		Avg	55	46.65	55	55	55	49.15	54.99	55	48.50
		Max	100	50.45	100	100	100	55.01	99.97	100	52
9	1.35	Min	0	57.14	0	54.26	0	49.92	55.97	0	54.33
		Avg	71.77	60.58	71.77	61.14	71.77	62	60.81	71.77	61.13
		Max	80	65	80	65	80	65	65	80	65
10	11.33	Min	21	28.99	0	21	0	20.34	21	0	21
		Avg	42.31	40.63	38.75	42.31	38.75	42.20	42.31	38.75	42.31
		Max	50	43.32	100	50	100	51.58	50	100	50
11	6.48	Min	22	35.34	20	20	20	20	35.34	20	20.95
		Avg	30.67	35.34	30.90	30.90	30.90	30.90	35.34	30.90	31.18
		Max	100	35.34	51	51	51	51	35.34	51	50.39
12	1.71	Min	15	37.37	15	23.57	15	33.57	15	15	33.44
		Avg	50.31	48.41	50.31	49.58	50.31	45.66	50.31	50.31	45.57
		Max	70	65.26	70	68.19	70	69.65	70	70	70.75

The BR strategies are unknown prior to playing the game because the regressed $\hat{\tau}$ -values require empirical estimation, implying that the BR model is best utilized as a benchmark rather than a prescriptive model in its own right. The BR model is designed to maximize a player's expected reward; therefore, when utilizing the average expected payoff as a comparative metric, it does generally yield the best results across all nine models tested. Although there are two exceptions (i.e., Games 2 and 11), we suspect these results may relate to the quality of fit of these specific

τ -estimates provided by Camerer et al. (2004)⁴.

The MM model yields the maximum of the minimum expected payoffs across all games, but it generally sacrifices performance with respect to both the average expected payoff and the maximum expected payoff. The conservativeness of the MM strategies can be observed by comparing the BR and MM columns across the average and maximum expected payoff results in Table 27.

Regarding the PE method, we note that the average expected payoff of its strategies typically do not exceed those of the BR strategies (with the exception of Game 11). This result aligns with the theory presented by Camerer et al. (2004). Likewise, for this particular set of games and empirical play, the PE strategies often yield both lower minimum and average expected payoffs compared to the MM strategies (i.e., Games 2, 4, 10, and 11).

The behaviorally robust strategies (i.e., R1 – DR2) provide compromise solutions between the conservative MM strategies and the unknown BR strategies. The behaviorally robust strategies trade improved average performance for decreased minimum performance relative to the MM strategies; however, their success in doing so depends upon the accuracy of the underlying information conditions. Consider Game 10 with $\hat{\tau} = 11.33$. Information conditions S1 and S2 both assume a high probability of τ existing in [1,2] and a low probability of larger τ -values. As with the PE strategy for this game, the S1 and S2 strategies are based upon an inadequate understanding of the underlying uncertainty and yield poor minimum and average results compared to the BR and MM strategies. Alternatively, for Game 4, each behaviorally robust model appropriately balances an improved average expected payoff with a reduced minimum

⁴For Game 2 with $\tau = 0$, we have $CH_{row} = (0.33, 0.33, 0.33)$ whereas the empirical results of Stahl and Wilson (1995) report $R_{SW2} = (0.625, 0.25, 0.125)$. Running the CH model for Game 11 with $\tau = 6.48$ and $k_{max} = 6$, the resulting solution is approximately equal to the empirical results, R_{SW11} , of Stahl and Wilson (1995). That is, $CH_{row} = (0.27, 0.07, 0.66) \approx R_{SW11} = (0.27, 0.08, 0.65)$. However, at such a low value of k in relation to τ , the CH solution has not yet converged. Increasing to $k_{max}=50$, we have $CH_{row} = (0.47, 0.19, 0.34)$ at approximate convergence.

expected outcome because $\hat{\tau} = 2.34$ is well represented in the information conditions. Finally, we also observe that at least one of the behaviorally robust strategies always weakly dominates the PE strategies for each game in terms of minimum, average, and maximum expected payoff. Of note, the DR2 strategies weakly dominate the PE strategies in terms of average expected payoff for nine of the twelve games and strictly dominate them for four of the twelve.

6.5 Conclusion

The bounded rationality exhibited by humans in one-shot, simultaneous-move games is a source of uncertainty complicating the decision for a player who confronts them. For such a player to maximize his or her utility, they must accurately assess their opponents' ability to think strategically. The CH model developed by Camerer et al. (2004) provides a framework for assessing such strategic ability, but the defining parameter τ must be identified via empirical estimation. As such, it is unlikely to be known precisely *ex ante*.

Within this research, we have leveraged robust optimization, stochastic programming, and distributionally robust optimization techniques to develop six mathematical programming formulations to collectively identify behaviorally robust strategies under varying types of uncertainty regarding τ . Solutions to these math programming formulations provide a player with a prescriptive strategy in accordance with their knowledge and beliefs regarding their opponents.

Moreover, we also developed a software package, BRMaximin, to identify behaviorally robust strategies under varying forms of uncertainty and utilized these tools to analyze the 12 games introduced by Stahl and Wilson (1995). The results illustrated the differing strategies obtained in accordance with a player's understanding of the uncertainty about τ . Likewise, we tested these strategies against the empirical

play from the human subjects in Stahl and Wilson (1995) to demonstrate how the accuracy of this knowledge affects a player's actual received payoffs.

Notwithstanding these results, there exists substantial opportunity for future research in this setting. One such opportunity involves revisiting the models presented herein by varying the underlying CH assumptions. For example, specific τ -values may be incorporated for each opponent population instead of the global opponent τ -value utilized herein. Additional avenues for future inquiry include the modification of our mathematical programs to other nonequilibrium structural models (e.g., Chong et al., 2016) or the automation of exact interval-based uncertainty solution methods via a computer algebra system. In pursuing such an agenda, researchers will be better able to advise players on their interactions against boundedly rational opponents and further the collective understanding of how to play normal form games against human opponents.

VII. Conclusions

In this research, we have developed multiple models that inform the offensive application of and defensive measures against persuasion in a national security setting, by way of completing five specific research goals in two related research threads.

In the first research thread pertaining to offensive and defensive behavioral influence models, an offensive modeling framework is created to identify how an entity optimally influences a populace to take a desired course of action, a defensive modeling framework is defined wherein a regulating entity takes action to bound the behavior of multiple adversaries simultaneously attempting to persuade a group of decisionmakers, and an offensive influence modeling framework under conditions of ambiguity is developed in accordance with historical information limitations. In the second research thread pertaining to behavioral and behaviorally robust approaches to deterrence, we demonstrate the alternative insights behavioral game theory generates for military operations planning, and we define behaviorally robust models for an agent to use in a normal form game under varying forms of uncertainty in order to inform deterrence policy decisions.

In accomplishing these goals, we provide military and civilian planners with frameworks to better conduct and counter political warfare, and with alternative tools to develop deterrence policy that addresses and exploits uncertain adversarial behavior.

Moreover, whereas the efficacy of influence in international strategic competition is well-known (Clarke, 2018; Heath, 2018; McClintock, 2018; Robinson et al., 2018), its effectiveness is difficult to quantify (Boot and Doran, 2013). By leveraging quantitative psychological theories, this research provides methods to mathematically model these operations and, in doing so, extends this scope of operations research applications beyond conventional warfare and prepares it for future use in the age of political warfare.

Appendix A. Expected Utility Influence Model

The GPP can be readily adapted to alternative assumptions of an expected utility maximizing decisionmaker. To illustrate this point, we model the persuasion problem previously referenced as follows. Note that the i, j suffixes have been dropped for notational convenience because there is only one decisionmaker and two prospects, with q being equivalent to prospect B. Likewise, we assume $\Delta = 0$.

$$\begin{aligned}
 & \max \Phi \\
 & \text{subject to } a_2 + a_3 \leq 400 \\
 & a_1 \leq 0.1 \\
 & EU(A) = 250 \\
 & EU(B) = (0.4 + a_1)(300 + a_2) + (0.6 - a_1)(a_3) \\
 & EU(B) - EU(A) - s^{pos} + s^{neg} = 0 \\
 & s^{pos} \leq M(1 - \Psi) \\
 & s^{neg} \leq M\Psi \\
 & Ms^{pos} \geq z^q \\
 & \Phi \leq z^q \\
 & a_1, a_2, a_3, s^{pos}, s^{neg} \geq 0 \\
 & \Phi, \Psi \in \{0, 1\}.
 \end{aligned}$$

Solving this model with BARON yields an optimal solution of $\Phi = 1$ and $(a_1, a_2, a_3) = (0.170, 197.714, 202.286)$. This solution is one of a set of alternative optimal solutions that may exist. If desired, additional constraints can be added to the formulation to preclude the identification of this optimal solution, and the perturbed instance resolved to identify an alternative optimal solution. Such a process can be iteratively repeated to explore the set of alternative optimal solutions.

Appendix B. IISE Proceedings: *Challenges and Solutions
with Exponentiation Constraints using Decision Variables
via the BARON Commercial Solver*

Challenges and Solutions with Exponentiation Constraints using Decision Variables via the BARON Commercial Solver

William N. Caballero, Alexander G. Kline, and Brian J. Lunday
Department of Operational Sciences
Air Force Institute of Technology, WPAFB, OH 45433

Abstract

We observe and explore a persistent issue encountered with the commercial global optimization solver BARON, wherein the solver falsely declares problem instances of a particular math programming formulation as infeasible. Problematic to BARON, the formulation contains constraints having exponentiation with decision variables in both the base and the exponent. We compare BARON's performance for this math programming problem against other commercial solvers, explore the potential cause of the false infeasible termination, and demonstrate how to mitigate this error by perturbing the formulation.

Keywords

Global Optimization, BARON, Nonlinear Programming

1. Introduction

The Branch and Reduce Optimization Navigator (BARON) is a preeminent global optimization solver that utilizes "constraint propagation, interval analysis, and duality in its reduce arsenal with advanced branch-and-bound optimization concepts" to find optimal solutions to non-convex mathematical programming problems [1]. Foundational work with respect to its nonlinear function relaxations, range reduction strategies, and branching methods was developed by Tawarmalani and Sahinidis [2]. With specific regard to fractional programming problems (i.e., functions which can be decomposed as the sum and products of univariate functions), the authors discuss methods to develop a convex relation to a factorable function, and the creation of a polyhedral outer-approximation to this relaxation. That is, the described methods serve to create a linear approximation of some factorable function by first creating a convex approximation which is, in turn, approximated linearly with hyperplanes. The iterative application of these techniques incorporated into a branch-and-bound framework with the node partitioning and fathoming rules set forth by Tawarmalani and Sahinidis [2] form the foundation of BARON's algorithmic procedure. The Sahinidis Group [3] provides a systematic comparison of BARON to four other leading global optimization solvers on 1740 test instances which shows BARON is able to solve instances much quicker than other commercial solvers. Extensive testing has also been conducted by Neumaier *et al.* [4] who concluded that BARON is the fastest and most robust global solver currently available.

For these reasons, BARON has earned its developers much recognition, including the 2004 INFORMS Computing Society Prize and the 2006 Beale-Orchard-Hays Prize from the Mathematical Optimization Society [5, 6]. It is available for use under a variety of algebraic modeling languages, including AIMMS, AMPL, and GAMS. Likewise, BARON has been utilized for supply chain design, integrated process water networks, scheduling, molecular design, manufacturing, and healthcare [7–11]. Its reputation as a reliable and effective global solver is such that BARON is commonly used as a benchmark for heuristics [12, 13].

Despite this success, Lastusilta *et al.* [14] suggest the solver AlphaECP outperforms BARON over the test instances available in MINLPLib. Likewise, Neumaier *et al.* [4] reference errors wherein BARON falsely reports an instance as infeasible. We expand upon these results by demonstrating a persistent issue encountered with BARON resulting in a false declaration of infeasibility for a class of problems which have a constraint with a decision variable exponentiated to another decision variable, henceforth referred to as a *power program*. In Section 2, we describe a specific power program formulation inducing the error, examine the potential cause, and provide programmatic perturbations to address it. Numerical tests in Section 3 illustrate the prevalence of these errors and the efficacy of our perturbations.

2. Rank Dependent Power Program

The presented power program is related to multiple individuals being offered many risky prospects with the objective to maximize the sum of perceived gain probabilities across all individuals. Concepts are borrowed from Cumulative Prospect Theory (CPT) in this formulation in terms of ranking outcomes and probability distortion, but the programs are meant to stand alone as an example and do not necessarily abide by the tenants of CPT [15]. Of note, we utilize neither the concept of a reference point nor the cumulative probability weighting function.

We consider a scenario wherein individuals in a set I are offered a variety of prospects. Each individual $i \in I$ is offered n_i prospects from a set J_i , each of which has m_{ij} outcomes, indexed on the set K_{ij} . We assume an external decision maker is able to alter the offered outcomes' raw values and probabilities by some constant amount through some binary persuasion action. In doing so, this decision maker wishes to maximize the sum of perceived probabilities of positive outcome values across all individuals and prospects. We continue by introducing the requisite parameters and decision variables.

Parameters

- \hat{x}_{ijk} : Baseline raw value for k -th outcome of prospect j for individual i before persuasion
- \hat{p}_{ijk} : Baseline probability of k -th outcome of prospect j for individual i before persuasion
- ε : Arbitrary sufficiently small positive real number
- M : Arbitrary sufficiently large real number
- \hat{f}_{ijk} : Persuasion effect on outcome k raw value for prospect j and individual i
- \hat{g}_{ijk} : Persuasion effect on outcome k probability for prospect j and individual i

Primary Decision Variables

- $T_{ijkk'}^+$: Equal to 1 if x_{ijk} is the $(m_{ij} - 1 + k')$ th greatest gain for $k = 1, \dots, m_{ij}$, and 0 otherwise; defined for all (i, j) combinations
- $T_{ijkk'}^-$: Equal to 1 if x_{ijk} is the (k') th greatest loss for $k = 1, \dots, m_{ij}$, and 0 otherwise; defined for all (i, j) combinations
- γ_i : Gain distortion coefficient for individual i after persuasion
- a_i : Binary variable indicating action taken against individual i . Equals one if action taken, zero otherwise.

Intermediate Decision Variables

- x_{ijk} : Gain/loss for individual i for k -th outcome of prospect j after persuasion
- p_{ijk} : Probability of k -th outcome of prospect j for individual i after persuasion
- $t_{ijk'}^+$: Ascending rank based list of x_{ijk} gains corresponding with mapping $T_{ijkk'}^+$
- $t_{ijk'}^-$: Ascending rank based list of x_{ijk} losses corresponding with mapping $T_{ijkk'}^-$
- $b_{ijk'}^+$: Corresponding probabilities for sorted $t_{ijk'}^+$ outcomes
- $b_{ijk'}^-$: Corresponding probabilities for sorted $t_{ijk'}^-$ outcomes
- $\pi_{ijk'}^+$: Distorted gain probability for i on k th outcome of prospect j after persuasion

Many of the intermediate decision variables could be eliminated and substituted with their explicit functional form in the constraints. However, they are maintained for tractability purposes and to facilitate the performance of GAMS/BARON in accordance with published documentation [16]. Using these sets, parameters, and decision variables, we define our Rank Dependent Power Program.

Rank Dependent Power Program - Base Model

$$\max \sum_{i \in I} \sum_{j \in J_i} \sum_{k' \in K_{ij}} \pi_{ijk'}^+ \quad (1a)$$

subject to

$$x_{ijk} = \hat{x}_{ijk} + a_i \hat{f}_{ijk}, \quad \forall i \in I, j \in J_i, k \in K_{ij} \quad (1b)$$

$$p_{ijk} = \hat{p}_{ijk} + a_i \hat{g}_{ijk}, \quad \forall i \in I, j \in J_i, k \neq m_{ij} \quad (1c)$$

$$p_{ijk} = 1 - \sum_{l=1}^{m_{ij}-1} p_{ijl}, \quad \forall i \in I, j \in J_i, k = m_{ij} \quad (1d)$$

$$\sum_{k=1}^{m_{ij}} p_{ijk} = 1, \quad \forall i \in I, j \in J_i \quad (1e)$$

$$\sum_{k=1}^{m_{ij}} T_{ijkk'}^+ + T_{ijkk'}^- = 1, \quad \forall i \in I, j \in J_i \quad (1f)$$

$$\sum_{k'=1}^{m_{ij}} T_{ijkk'}^+ + T_{ijkk'}^- = 1, \quad \forall i \in I, j \in J_i \quad (1g)$$

$$t_{ijk'}^+ = \sum_{k=1}^{m_{ij}} T_{ijkk'}^+ x_{ijk}, \quad \forall i \in I, j \in J_i, k' \in K_{ij} \quad (1h)$$

$$t_{ijk'}^- = \sum_{k=1}^{m_{ij}} T_{ijkk'}^- x_{ijk}, \quad \forall i \in I, j \in J_i, k' \in K_{ij} \quad (1i)$$

$$t_{ij(k'+1)}^- \geq t_{ijk'}^-, \quad \forall i \in I, j \in J_i, k = 1, \dots, m_{ij} - 1 \quad (1j)$$

$$t_{ij(k'+1)}^+ \geq t_{ijk'}^+, \quad \forall i \in I, j \in J_i, k = 1, \dots, m_{ij} - 1 \quad (1k)$$

$$t_{ijk'}^+ \geq 0, \quad \forall i \in I, j \in J_i, k' \in K_{ij} \quad (1l)$$

$$t_{ijk'}^- \leq 0, \quad \forall i \in I, j \in J_i, k' \in K_{ij} \quad (1m)$$

$$b_{ijk'}^+ = \sum_{k=1}^{m_{ij}} T_{ijkk'}^+ p_{ijk}, \quad \forall i \in I, j \in J_i, k' \in K_{ij} \quad (1n)$$

$$b_{ijk'}^- = \sum_{k=1}^{m_{ij}} T_{ijkk'}^- p_{ijk}, \quad \forall i \in I, j \in J_i, k' \in K_{ij} \quad (1o)$$

$$\pi_{ijk'}^+ = (b_{ijk'}^+)^{\gamma_i}, \quad \forall i \in I, j \in J_i, k' \in K_{ij} \quad (1p)$$

$$a_i \in \{0, 1\}, \gamma_i \geq 1, \quad \forall i \in I$$

$$T_{ijkk'}^+, T_{ijkk'}^- \in \{0, 1\}, \quad \forall i \in I, j \in J_i, k, k' \in K_{ij}$$

The objective function (1a) maximizes the sum of distorted probabilities for gains across all individuals, prospects, and outcomes. Constraints (1b) through (1e) serve to update the offered outcomes and their probabilities. Constraints (1f) through (1o) ensure all prospects are sorted in ascending order and labeled as a gain or a loss depending on the values of $T_{ijkk'}^+$ and $T_{ijkk'}^-$. Constraints (1f) through (1m) create the mappings, $T_{ijkk'}^+$ and $T_{ijkk'}^-$. Constraints (1f) and (1g) enforce a bijective mapping. Constraints (1h) and (1i) perform the actual mapping calculation for outcomes, and Constraints (1n) and (1o) do likewise for probabilities. Constraints (1j) and (1k) ensure the values are in ascending order, and Constraints (1l) and (1m) enforce the positivity and negativity of gains and losses, respectively. Constraint (1p) serves as an intermediate decision variable of the terms summed in the objective function. This equality is responsible for converting the formulation into a power program.

For a given set of a_i decision variables, Constraints (1b) – (1p) are completely determined. This is readily noted by observing that the decision variables a_i dictate the values of x_{ijk} and p_{ijk} , which in turn control all remaining decision variables except γ_i . The values $T_{ijkk'}^+$ and $T_{ijkk'}^-$ are easily found by sorting the x_{ijk} -values and observing their signs. The remaining intermediate decision variables can then be found by simple calculation. Furthermore, this formulation always has a feasible solution, regardless of parameter values. This can be observed from the following: (1) Constraints (1b), (1c), (1d), and (1p) assign a value to free decision variables, and (2) A subset of these assignments is used in Constraints (1f)–(1o) to assign more free variables or sort them in ascending order of their associated outcome.

Using these results and observing that increasing γ_i always decreases the objective function value based on the domain of b_{ijk}^+ , finding an optimal solution to this program can be simplified to searching among all possible combinations of a_i decision variables. This combinatorial construct is not necessarily helpfully for large instances, but allows us to easily find optimal solutions to small instances of the program for testing purposes.

BARON is unable to directly handle a function x^y when both x and y are decision variables [17]. In order to process such functions, GAMS/BARON transforms the function to $e^{y \log(x)}$. However, the domains of the two functions are not equivalent. For example, the original function is defined at $x = 0$, whereas the transformation is not.

In an effort to resolve this conflict, we introduce the reformulations A1 through A4. Each perturbation is based on the hypothesis that pre-processing bounds for p_{ijk} are not communicated to b_{ijk}^+ and an error is triggered when BARON observes the possibility of taking $\log(0)$. A1 removes the intermediate decision variable, against the general guidance of GAMS documentation. However, by removing the intermediate decision variable, any error associated with its bounds should be eliminated. A2 and A3 attempt to add a small real number to the intermediate decision variables in (1p) to avoid a domain violation in the transformation. A4 provides a similar solution by lower bounding the sorted probabilities by some very small real number. That is, b_{ijk}^+ is explicitly stated as greater than zero. Strictly speaking, such a lower bound makes the program infeasible. However, if chosen small enough (e.g., 1×10^{-16}), the lower bound is essentially treated as a roundoff error while remaining defined in the transformed power function.

Alternative or Additional Constraints

$$A1 : \pi_{ijk}^+ = \sum_{k=1}^{m_{ij}} T_{ijkk}^- p_{ijk}^{\gamma_i} \text{ substituted for (1p)}$$

$$A2 : \pi_{ijk}^+ = (b_{ijk}^+ + \epsilon)^{\gamma_i} \text{ substituted for (1p)}$$

$$A3 : \pi_{ijk}^+ = |b_{ijk}^+ + \epsilon|^{\gamma_i} \text{ substituted for (1p)}$$

$$A4 : b_{ijk}^+ \geq \epsilon, b_{ijk}^- \geq \epsilon \text{ added to the base model}$$

3. Testing and Analysis

We compare BARON's performance on 100 instances of the base Rank Dependent Power Program with two individuals each being offered two prospects, for which there respectively exist two potential outcomes, to that of five other solvers. Of these five solvers, one is a global solver (i.e., SCIP) while the other four (i.e., DICOPT, LINDO, SSB, and AlphaECP) are primarily used for solving convex MINLPs. The 100 random instances draw parameters from the following distributions such that $x_{ijk} \in [-6500, 6500]$, $p_{ijk} \in [0.01, 0.99]$, and $e^{\gamma_i \log(b_{ijk}^+)}$ is always defined: $\hat{x}_{ijk} \sim U[-5000, 5000]$, $\hat{p}_{ijk} \sim U[0.25, 0.75]$, $\hat{f}_{ijk} \sim U[-1500, 1500]$, and $\hat{g}_{ijk} \sim U[-0.249, 0.249]$. To calculate the optimality gap, the action space in each instance is enumerated and its effect on (1b) – (1p) manually calculated as previously discussed.

Testing is performed on an HP ZBook equipped with a 2.70 GHz Intel i7-4800MQ processor and 32GB of RAM. Each solver is provided a relative optimality termination criteria of 0.001, an iteration limit of 3000, and a time limit of five minutes. Table 1 details the termination criterion invoked for each solver over the 100 random instances. Only AlphaECP correctly identifies all instances as feasible. The global solvers BARON and SCIP struggle to find feasible solutions and terminate with a conclusion of infeasibility for 100 and 85 instances, respectively. DICOPT does not label any of the instances as infeasible but is only able to return a solution for 16 of the instances. SSB and LINDO incorrectly terminate by designating 8 and 12 instances as infeasible, respectively. Taken collectively, the global solvers appear to underperform their convex MINLP counterparts.

With regard to BARON, we compare the alternative formulations and examine the termination criterion. It can be observed in Table 1 that A1, A3, and A4 all conclude that 95 instances are feasible. Surprisingly, alternative A2 concludes that only 2 instances are feasible. However, upon examination, the evaluation of the objective function is found to be incorrect. These two observations are excluded from further analysis. We postulate the difference in performance between A2 and A3 derives from the pre-processing technique utilized for the absolute value function.

Although the alternative formulations A1, A3, and A4 greatly improve performance, erroneous infeasible terminations still occur. AlphaECP is the only solver which correctly identifies the feasibility of all instances. However, Table 1 shows that, for the instances BARON identifies as feasible, it has the lowest average optimality gap.

Table 1: Feasibility and Optimality Comparisons

Solver	# Feasible	# Infeasible	# Timed Out	Avg RelGap	Avg AbsGap
BARON (Base Model)	0	100	0	-	-
BARON (A1)	95	5	0	0.59%	0.0141
BARON (A2)	2	98	0	-	-
BARON (A3)	95	5	0	0.42%	0.0102
BARON (A4)	95	5	0	6.72%	0.1611
SCIP (Base Model)	15	85	0	36.68%	0.5227
DICOPT (Base Model)	16	0	84	37.18%	0.5525
AlphaECP (Base Model)	100	0	0	17.74%	0.3097
SSB (Base Model)	92	8	0	14.64%	0.3236
LINDO (Base Model)	84	12	4	17.68%	0.3798

4. Conclusions

We have examined a problem which belongs to the very challenging class of power programs. The program analyzed proves to be problematic for the GAMS/BARON solver combination. In its base form, an incorrect infeasible termination is reached in the Rank Dependent Power Program for each of the hundred instances examined which is the worst result over all solvers considered. These results are believed to be due to the GAMS/BARON transformation utilized for x^y when both x, y are decision variables. Accounting for this transformation, alternative formulations are examined which yield promising results in terms of solution quality, but still conclude with incorrect infeasible terminations for some instances.

Although, BARON is advertised as not requiring an initial seeded solution, for problems such as ours, it may be a helpful tactic to mitigate the observed error. Since AlphaECP seemingly is the most effective software in terms of finding a feasible solution, a combined approach of using AlphaECP to find a starting point which is then fed into BARON may prove efficacious.

Fortunately, there always exists a feasible solution for the examined power program in this research. However, our results demonstrate the potential for BARON to yield a false infeasible termination, which is problematic for instances without a guarantee of feasibility. In order to mitigate this issue, we have provided a collection of programmatic mitigation techniques for practitioner use.

References

1. The Optimization Firm, 2017, "BARON", <http://www.minlp.com/baron>, Accessed November 7, 2017.
2. Tawarmalani, Mohi, and Sahinidis, Nikolaos V., 2004, "Global Optimization of Mixed-Integer Nonlinear Programs: A Theoretical and Computational Study", *Mathematical Programming*, 99(3), 563-591.
3. Sahinidis Optimization Group, 2017, "BARON Software: Comparisons", <http://archimedes.cheme.cmu.edu/?q=baron>, Accessed August 31, 2017.
4. Neumaier, Arnold, Shcherbina, Oleg, Huyer, Waltraud, and Vinko, Tamas, 2005, "A Comparison of Complete Global Optimization Solvers", *Mathematical Programming*, 103(2), 335-356.
5. Institute for Operations Research and Management Science, 2017, "INFORMS Computing Society (ICS) Prize", <http://connect.informs.org/computing/awards/ics-prize>, Accessed October 1, 2017.
6. Mathematical Programming Society, 2006, "2006 Beale-Orchard-Hays Prize Citation", http://www.mathprog.org/prz/citations/boh_2006.htm, Accessed October 1, 2017.
7. Ahmetovic, Elvis, and Grossmann, Ignacio E., 2011, "Global Superstructure Optimization for the Design of Integrated Process Water Networks", *AIChE Journal*, 57(2), 434-457.
8. Ahmetovic, Elvis, Martin, Mariano, and Grossmann, Ignacio E., 2010, "Optimization of Energy and Water Consumption in Corn-based Ethanol Plants", *Industrial & Engineering Chemistry Research*, 49(17), 7972-7982.
9. Mahar, Stephen, Bretthauer, Kurt M., and Salzarulo, Peter A., 2011, "Locating Specialized Service Capacity in a Multihospital Network", *European Journal of Operational Research*, 212(3), 596-605.
10. Sahinidis, Nikolaos V., and Tawarmalani, Mohit, 2000, "Applications of Global Optimization to Process and Molecular Design", *Computers & Chemical Engineering*, 24(9), 2157-2169.
11. You, Fengqi, and Grossmann, Ignacio E., 2008, "Design of Responsive Supply Chains Under Demand Uncertainty", *Computers & Chemical Engineering*, 32(12), 3090-3111.

12. Lunday, Brian J., and Sherali, Hanif D., 2012, "Minimizing the Maximum Network Flow: Models and Algorithms with Resource Synergy Considerations", *Journal of the Operational Research Society*, 63(12), 1693-1707.
13. Nowak, Ivo, and Vigerske, Stefan, 2008, "LaGO: A (heuristic) Branch and Cut Algorithm for Nonconvex MINLPs". *Central European Journal of Operations Research*, 16(2), 127-138.
14. Lastusilta, Toni, Bussieck, Michael R., and Westerlund, Tapio, 2009, "An Experimental Study of the GAMS/AlphaECP MINLP Solver", *Industrial & Engineering Chemistry Research*, 48(15), 7337-7345.
15. Tversky, Amos, and Kahneman, Daniel, 1992, *Advances in Prospect Theory: Cumulative Representation of Uncertainty*. *Journal of Risk and Uncertainty*, 5(4), 297-323, 1992.
16. GAMS, 2017, "GAMS Documentation 24.9: Good NLP Formulations", https://www.gams.com/latest/docs/UG_NLP_GoodFormulations.html, Accessed Aug 25, 2017
17. Sahinidis, Nikolaos V., 2017, "BARON User Manual v. 17.8.9", <http://www.minlp.com/downloads/docs/baron%20manual.pdf>, Accessed Aug 26, 2017 .

Bibliography

- Abdellaoui, M. (2000). Parameter-Free Elicitation of Utility and Probability Weighting Functions. *Management Science*, 46(11):1497–1512.
- Abdellaoui, M., Bleichrodt, H., and Paraschiv, C. (2007). Loss Aversion under Prospect Theory: A Parameter-free Measurement. *Management Science*, 53(10):1659–1674.
- Aghassi, M. and Bertsimas, D. (2006). Robust Game Theory. *Mathematical Programming*, 107:231–273.
- Air University (2018). Influence Operations: Definitions. <http://www.au.af.mil/info-ops/index.htm>. Retrieved on 20 Apr 2018.
- Ajzen, I. (1991). The Theory of Planned Behavior. *Organizational Behavior and Human Decision Processes*, 50(2):179–211.
- Andonie, C. and Diermeier, D. (2017). Path-dependency and Coordination in Multi-candidate Elections with Behavioral Voters. *Journal of Theoretical Politics*, 29(4):520–545.
- Andreoni, J. and Miller, J. H. (1993). Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma: Experimental Evidence. *The Economic Journal*, 103(May):570 – 585.
- Andreoni, J. and Sprenger, C. (2010). Certain and Uncertain Utility: The Allais Paradox and Five Decision Theory Phenomena. Levine’s working paper archive, David K. Levine.
- Apolte, T. (2019). I Hope I Die Before I Get Old. *Defence and Peace Economics*, Advance online publication.
- Ares, G. and Varela, P. (2018). *Methods in Consumer Research, Volume 2: Alternative Approaches and Special Applications*. Woodhead Publishing, Cambridge, MA.
- Argyris, N. and French, S. (2017). Nuclear Emergency Decision Support: A Behavioural OR Perspective. *European Journal of Operational Research*, 262(1):180–193.
- Ariely, D. and Loewenstein, G. (2006). The Heat of the Moment: The Effect of Sexual Arousal on Sexual Decision Making. *Journal of Behavioral Decision Making*, 19(2):87–98.
- Athans, M. (1987). Command and Control (C2): A Challenge to Control Science. *IEEE Transactions on Automatic Control*, 32(4):286–293.

- Audet, C. and Dennis Jr., J. E. (2006). Mesh Adaptive Direct Search Algorithms for Constrained Optimization. *SIAM Journal on Optimization*, 17(1):188–217.
- Aumann, R. and Brandenburger, A. (1995). Epistemic Conditions for Nash Equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1161–1180.
- Australian Department of Defence (2009). Australian Defence Doctrine Publication 5.0: Joint Planning.
- Ba, S., Myers, W. R., and Brenneman, W. A. (2015). Optimal Sliced Latin Hypercube Designs. *Technometrics*, 57(4):479–487.
- Baliga, S. and Sjöström, T. (2004). Arms Races and Negotiations. *The Review of Economic Studies*, 71(2):351–369.
- Barlow, H. (1961). Possible Principles Underlying the Transformations of Sensory Messages. In Rosenblith, W., editor, *Sensory Communication*, pages 217–234. MIT Press, Cambridge, MA.
- Becker, K. H. (2016). An Outlook on Behavioural OR—Three Tasks, Three Pitfalls, One Definition. *European Journal of Operational Research*, 249(3):806–815.
- Becker-Peth, M. and Thonemann, U. W. (2016). Reference Points in Revenue Sharing Contracts - How to Design Optimal Supply Chain Contracts. *European Journal of Operational Research*, 249(3):1033–1049.
- Belk, R. W. (2007). *Handbook of Qualitative Research Methods in Marketing*. Edward Elgar Publishing, Northampton, MA.
- Bertsimas, D., B, B. D., and Caramanis, C. (2011). Theory and Applications of Robust Optimization. *SIAM Review*, 53(3):464–501.
- Bhatt, M. A. and Camerer, C. F. (2005). Self-referential Thinking and Equilibrium as States of Mind in Games: fMRI Evidence. *Games and Economic Behavior*, 52(2):424–459.
- Bhatt, M. A., Lohrens, T., Camerer, C. F., and Montague, P. R. (2010). Neural Signatures of Strategic Types in a Two-person Bargaining Game. *Proceedings of the National Academy of Sciences*, 107(46):1–6.
- Boardman, N. T., Lunday, B. J., and Robbins, M. J. (2017). Heterogeneous Surface-to-Air Missile Defense Battery Location: A Game Theoretic Approach. *Journal of Heuristics*, 23(6):417–447.
- Booij, A. S. and Van de Kuilen, G. (2009). A Parameter-free Analysis of the Utility of Money for the General Population under Prospect Theory. *Journal of Economic Psychology*, 30(4):651–666.

- Booij, A. S., Van Praag, B. M., and Van De Kuilen, G. (2010). A Parametric Analysis of Prospect Theory's Functionals for the General Population. *Theory and Decision*, 68(1-2):115–148.
- Boot, M. and Doran, M. (2013). Political Warfare. *Council on Foreign Relations: Policy Innovation Memorandum*.
- Boyd, C. (2011). The Future of MISO. *Special Warfare*, 24(1):22–28.
- Brader, T. (2005). Striking a Responsive Chord: How Political Ads Motivate and Persuade Voters by Appealing to Emotions. *American Journal of Political Science*, 49(2):388–405.
- Brandt, U. S. and Svendsen, G. T. (2018). How Robust is the Welfare State When Facing Open Borders? An Evolutionary Game-theoretic Model. *Public Choice*, pages 1–17.
- Brown, M., Haskell, W. B., and Tambe, M. (2014). Addressing Scalability and Robustness in Security Games with Multiple Boundedly Rational Adversaries. In *Proceedings of GameSec 2014*, pages 23–42.
- Burke, J. V., Curtis, F. E., Lewis, A. S., Overton, M. L., and Simões, L. E. (2018). Gradient Sampling Methods for Nonsmooth Optimization. *arXiv preprint arXiv:1804.11003*.
- Burne, A. H. (2016). *The Crecy War: A Military History of the Hundred Years War from 1337 to the Peace of Bretigny in 1360*. Frontline Books, Barnsley, South Yorkshire.
- Byrd, W. A. (2017). Disease or Symptom? Afghanistan's Burgeoning Opium Economy in 2017. Afghanistan Research and Evaluation Unit.
- Caballero, W. N., Lunday, B. L., and Kline, A. G. (2018). Challenges and Solutions with Exponentiation Constraints using Decision Variables via the BARON Commercial Solver. In Barker, K., Berry, D., and Rainwater, C., editors, *2018 IISE Annual Conference Proceedings*, pages 1331–1336.
- Cai, Y., Candogan, O., Daskalakis, C., and Papadimitriou, C. (2016). Zero-Sum Polymatrix Games: A Generalization of Minmax. *Mathematics of Operations Research*, 41(2):648–655.
- Camerer, C., Loewenstein, G., and Prelec, D. (2005). Neuroeconomics: How Neuroscience Can Inform Economics. *Journal of Economics Literature*, 43(1):9–64.
- Camerer, C. F. (2000). *Prospect Theory in the Wild: Evidence from the Field*. California Institute of Technology, Pasadena, CA.

- Camerer, C. F. (2011). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, NJ.
- Camerer, C. F. and Ho, T.-H. (1998). Experience-weighted Attraction Learning in Coordination Games: Probability Rules, Heterogeneity, and Time-variation. *Journal of Mathematical Psychology*, 42:305–326.
- Camerer, C. F. and Ho, T.-H. (1999). Experience-weighted Attraction Learning in Normal Form Games. *Econometrica*, 67(4):827–874.
- Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2002). Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games. *Journal of Economic Theory*, 104(1):137–188.
- Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, 119(3):861–898.
- Campos-Vazquez, R. M. and Cuijy, E. (2014). The Role of Emotions on Risk Aversion: A Prospect Theory Experiment. *Journal of Behavioral and Experimental Economics*, 50:1–9.
- Caplan, B. (2006). Terrorism: The Relevance of the Rational Choice Model. *Public Choice*, 128(1-2):91–107.
- Castañeda, A. and Martinelli, C. (2018). Politics, Entertainment and Business: A Multisided Model of Media. *Public Choice*, 174(3-4):239–256.
- Cavagnaro, D., Pitt, M., Gonzalez, R., and Myung, J. (2013). Discriminating Among Probability Weighting Functions using Adaptive Design Optimization. *Journal of Risk and Uncertainty*, 47(3):255–289.
- Central Statistics Organization of Afghanistan (2007). Badakhshan: A Socio-economic and Demographic Profile. CSO and UNFPA.
- Central Statistics Organization of Afghanistan (2013). Afghanistan Living Conditions Survey 2011-2012. CSO.
- Central Statistics Organization of Afghanistan (2018). Afghanistan Living Conditions Survey 2016-2017. CSO.
- Cetin, E. and Esen, S. T. (2006). A Weapon–Target Assignment Approach to Media Allocation. *Applied Mathematics and Computation*, 175(2):1266–1275.
- Chakraborty, A. and Harbaugh, R. (2010). Persuasion by Cheap Talk. *The American Economic Review*, 100(5):2361–2382.
- Chiozza, G. (2017). Presidents on the Cycle: Elections, Audience Costs, and Coercive Diplomacy. *Conflict Management and Peace Science*, 34(1):3–26.

- Chong, J.-K., Ho, T.-H., and Camerer, C. (2016). A Generalized Cognitive Hierarchy Model of Games. *Games and Economic Behavior*, 99:257–274.
- Chughtai, A. (2018). Afghanistan: Who Controls What. <https://www.aljazeera.com/indepth/interactive/2016/08/afghanistan-controls-160823083528213.html>.
- Clarke, C. (2018). How Hezbollah Came to Dominate Information Warfare. <https://www.rand.org/blog/2017/09/how-hezbollah-came-to-dominate-information-warfare.html>. Retrieved on 18 Apr 2018.
- Clausewitz, C. v. (1989). *On War*, volume 1. Princeton University Press, Princeton, NJ.
- Clemens, J. (2008). Opium in Afghanistan: Prospects for the Success of Source Country Drug Control Policies. *The Journal of Law and Economics*, 51(3):407–432.
- Cohen, R. and Robinson, L. (2018). Political Warfare Is Back with a Vengeance. RAND Editorials. Retrieved on 18 Apr 2018.
- Coll, S. (2019). *Directorate S: The CIA and America's Secret Wars in Afghanistan and Pakistan*. Penguin Books, New York, NY.
- Colson, B., Marcotte, P., and Savard, G. (2007). An Overview of Bilevel Optimization. *Annals of Operations Research*, 153(1):235–256.
- Coricelli, G. and Nagel, R. (2009). Neural Correlates of Depth of Strategic Reasoning in Medial Prefrontal Cortex. *Proceedings of the National Academy of Sciences*, 106(23).
- Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2001). Cognition and Behavior in Normal-form Games: An Experimental Study. *Econometrica*, 69(5):1193–1235.
- Crawford, V. P. (2003). Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions. *American Economic Review*, 93(1):133–149.
- Crawford, V. P. and Sobel, J. (1982). Strategic Information Transmission. *Econometrica: Journal of the Econometric Society*, 50(6):1431–1451.
- Crockett, M. J. and Fehr, E. (2014). Pharmacology of Economic and Social Decision Making. In Glimcher, P. W. and Fehr, E., editors, *Neuroeconomics: Decision Making and the Brain*, chapter 14, pages 259–279. Elsevier, San Diego, CA.
- Croson, R., Boles, T., and Murnighan, J. K. (2003). Cheap Talk in Bargaining Experiments: Lying and Threats in Ultimatum Games. *Journal of Economic Behavior & Organization*, 51(2):143–159.

- Cruz, J. B., Simaan, M. A., Gacic, A., Jiang, H., Letellier, B., Li, M., and Liu, Y. (2001). Game-theoretic Modeling and Control of a Military Air Operation. *IEEE Transactions on Aerospace and Electronic Systems*, 37(4):1393–1405.
- Davidson, J. R., Connor, K. M., and Swartz, M. (2006). Mental Illness In U.S. Presidents Between 1776 and 1974: A Review of Biographical Sources. *Journal of Nervous and Mental Disease*, 194(1):47 – 51.
- Dempe, S. (2002). *Foundations of Bilevel Programming*. Springer Science & Business Media, New York, NY.
- DHS and FBI (2016). JAR-16-20296A: GRIZZLY STEPPE Russian Malicious Cyber Activity. https://www.us-cert.gov/sites/default/files/publications/JAR_16-20296A_GRIZZLY%20STEPPE-2016-1229.pdf. Retrieved on 17 Oct 2017.
- Diamond, S. S. (2003). Truth, Justice, and the Jury. *Harvard Journal of Law and Public Policy*, 26:143.
- Dupuy, B. (2018). North Korea Could Join Russia in Interfering in Next U.S. Election, Senators Warn. <https://www.newsweek.com/senators-warn-north-korea-could-target-elections-782316>. Retrieved on 9 August 2018.
- Eisenegger, C., Knoch, D., Ebstein, R. P., Gianotti, L. R., Sándor, P. S., and Fehr, E. (2010). Dopamine Receptor D4 Polymorphism Predicts the Effect of L-DOPA on Gambling Behavior. *Biological Psychiatry*, 67(8):702–706.
- Fairchild, R. (2014). Emotions in the Financial Markets. In Baker, H. K. and Ricciardi, V., editors, *Investor Behavior*, pages 347–364. John Wiley & Sons, Inc., Hoboken, New Jersey.
- Fearon, J. D. (1994). Domestic Political Audiences and the Escalation of International Disputes. *American Political Science Review*, 88(3):577–592.
- Fearon, J. D. (2011). Arming and Arms Races . In *Annual Meetings of the American Political Science Association, Washington, DC*.
- Fearon, J. D. (2018). Cooperation, Conflict, and the Costs of Anarchy. *International Organization*, 72(3):523–559.
- Fedeli, S., Leonida, L., and Santoni, M. (2018). Bureaucratic institutional design: the case of the Italian NHS. *Public Choice*, 177(3-4):265–285.
- Felbab-Brown, V. (2016). High and Low Politics in Afghanistan: The Terrorism-Drugs Nexus and What Can be Done About It. <https://www.brookings.edu/articles/high-and-low-politics-in-afghanistan-the-terrorism-drugs-nexus-and-what-can-be-done-about-it/n>. Brookings Institute.

- Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford University Press, Palo Alto, CA.
- Financial Services Authority (2018). Statutory Objectives. <http://www.fsa.gov.uk/about/aims/statutory>. Accessed 5 Jan 2019.
- Fishstein, P. (2014). Evolving Terrain: Opium Poppy Cultivation in Balkh and Badakhshan Provinces in 2013. Afghanistan Research and Evaluation Unit.
- Fox, C. R. (1999). Strength of Evidence, Judged Probability, and Choice under Uncertainty. *Cognitive Psychology*, 38(1):167–189.
- Fox, C. R. and Tversky, A. (1998). A Belief-based Account of Decision Under Uncertainty. *Management Science*, 44(7):879–895.
- Franco, L. A. and Hämäläinen, R. P. (2016). Engaging with Behavioural OR: On Methods, Actors, and Praxis. In *Behavioural Operational Research: Theory, Methodology and Practice*, pages 3–26. Palgrave Macmillan, London, UK.
- Fry, J. and Binner, J. M. (2016). Elementary Modelling and Behavioural Analysis for Emergency Evacuations Using Social Media. *European Journal of Operational Research*, 249(3):1014–1023.
- Gabrel, V., Murat, C., and Theile, A. (2014). Recent Advances in Robust Optimization: An Overview. *European Journal of Operational Research*, 235(3):471–483.
- Gächter, S., Johnson, E. J., and Herrmann, A. (2007). Individual-level Loss Aversion in Riskless and Risky Choices. Technical report, University of Nottingham.
- Gass, R. H. and Seiter, J. S. (2015). *Persuasion: Social Influence and Compliance Gaining*. Routledge, New York, NY.
- Georganas, S., Healy, P. J., and Weber, R. A. (2015). On the Persistence of Strategic Sophistication. *Journal of Economic Theory*, 159:369–400.
- Gigerenzer, G. and Selten, R. (2002). *Bounded Rationality: The Adaptive Toolbox*. MIT Press, Cambridge, MA.
- Gladwell, M. (2008). *Outliers: The Story of Success*. Hachette Book Group, New York, NY.
- Goddard, B. (2010). Cold Warriors Say No Nukes. <http://thehill.com/opinion/columnists/ben-goddard/78391-cold-warriors-say-no-nukes>. Retrieved on 14 June 2018.
- Goerigk, M. and Schöbel, A. (2016). Algorithm Engineering in Robust Optimization. In *Algorithm Engineering*, pages 245–279. Springer.

- Golman, R. and Page, S. E. (2009). General Blotto: Games of Allocative Strategic Mismatch. *Public Choice*, 138(3):279–299.
- Gonzalez, R. and Wu, G. (1999). On the Shape of the Probability Weighting Function. *Cognitive Psychology*, 38(1):129–166.
- Hafner-Burton, E., Haggard, S., Lake, D. A., and Victor, D. G. (2017). The Behavioral Revolution and International Relations. *International Organization*, 71(S1):1 – 31.
- Hämäläinen, R. P., Luoma, J., and Saarinen, E. (2013). On the Importance of Behavioral Operational Research: The Case of Understanding and Communicating about Dynamic Systems. *European Journal of Operational Research*, 228(3):623–634.
- Hanley Jr., J. T. (2017a). Changing DoD’s Analysis Paradigm. *Naval War College Review*, 70(1):64–103.
- Hanley Jr., J. T. (2017b). Planning for the Kamikazes: Toward a Theory and Practice of Repeated Operational Games. *Naval War College Review*, 70(2):29–48.
- Harsanyi, J. C. (1967). Games with Incomplete Information Played by Bayesian Players Parts I-III. *Management Sciences*, 14(3):159–182.
- Haywood, O. (1954). Military Decision and Game Theory. *Operations Research*, 2(4):365–385.
- Heath, T. (2018). Beijing’s Influence Operations Target Chinese Diaspora. <https://www.rand.org/blog/2018/02/beijings-influence-operations-target-chinese-diaspora.html>. Retrieved on 18 Apr 2018.
- Henrich, J. (200). Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining among the Machiguenga of the Peruvian Amazon. *American Economic Review*, 90(4):973 – 979.
- Herald, S. (2016). Four Changes, Four Courses Eyed for ‘Blended’ Retirement . <http://www.sunherald.com/news/local/military/article61591977.html>. Retrieved on 14 Feb 2018.
- Hertwig, R., Herzog, S. M., Schooler, L. J., and Reimer, T. (2008). Fluency Heuristic: A Model of How the Mind Exploits a By-product of Information Retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(5):1191.
- Ho, T. H., Camerer, C. F., and Chong, J. K. (2007). Self-tuning Experience Weighted Attraction Learning in Games. *Journal of Economic Theory*, 133(1):177–198.
- Homel, P. and Carroll, T. (2009). Moving Knowledge into Action: Applying Social Marketing Principles to Crime Prevention. *Trends & Issues in Crime and Criminal Justice*, 381(1):1–6.

- Horowitz, M. C. and Fuhrmann, M. (2018). Studying Leaders and Military Conflict: Conceptual Framework and Research Agenda. *Journal of Conflict Resolution*, 62(10):2072–2086.
- Imamura, A., Uitti, R. J., and Wszolek, Z. K. (2006). Dopamine Agonist Therapy for Parkinson Disease and Pathological Gambling. *Parkinsonism & Related Disorders*, 12(8):506–508.
- Johnson, D. D. and Fowler, J. H. (2011). The Evolution of Overconfidence. *Nature*, 477(7364):317.
- Jowett, G. S. and O'Donnell, V. (2014). *Propaganda & Persuasion*. SAGE Publications Sage CA: Los Angeles, CA.
- Kahn, H. (1960). *On Thermonuclear War*. Princeton University Press, Princeton, NJ.
- Kahneman, D. (2011). *Thinking Fast and Slow*. Farrar, Straus and Giroux, New York, NY.
- Kahneman, D. (2016). Heuristics and Biases. In *Scientists Making a Difference: One Hundred Eminent Behavioral and Brain Scientists Talk about Their Most Important Contributions*, pages 171–174. Cambridge University Press.
- Kahneman, D. and Tversky, A. (1972). Subjective Probability: A Judgment of Representativeness. *Cognitive Psychology*, 3(3):430–454.
- Kahneman, D. and Tversky, A. (1973). On the Psychology of Prediction. *Psychological Review*, 80(4):237–251.
- Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2):263–292.
- Kahneman, D. and Tversky, A. (1981). The Simulation Heuristic. Technical Report AD-A099504, Stanford University, Department of Psychology.
- Kahneman, D. and Tversky, A. (1982). Variants of Uncertainty. *Cognition*, 11(2):143–157.
- Kamenica, E. and Gentzkow, M. (2011). Bayesian Persuasion. *American Economic Review*, 101(6):2590–2615.
- Keeney, R. L. and Keeney, R. L. (2009). *Value-focused Thinking: A Path to Creative Decisionmaking*. Harvard University Press, Cambridge, MA.
- Keller, N. and Katsikopoulos, K. V. (2016). On the Role of Psychological Heuristics in Operational Research; and a Demonstration in Military Stability Operations. *European Journal of Operational Research*, 249(3):1063–1073.

- Kennan, G. (1948). Policy Planning Staff Memorandum. Technical Report Records of the National Security Council, NSC 10/2, United States Department of State.
- Kertzer, J. D., Rathburn, B., and Rathburn, N. (2018). The Price of Peace: A Behavioral Approach to Costly Signaling in International Relations. http://people.fas.harvard.edu/~jkertzer/Research_files/Price%20of%20Peace%20Web.pdf.
- Khajavirad, A. and Sahinidis, N. V. (2013). Convex Envelopes Generated from Finitely Many Compact Convex Sets. *Mathematical Programming*, 137(1–2):371–408.
- Kilka, M. and Weber, M. (2001). What Determines the Shape of the Probability Weighting Function under Uncertainty? *Management Science*, 47(12):1712–1726.
- Kofman, M. and Rojansky, M. (2015). *A Closer Look at Russia's "Hybrid War"*. Woodrow Wilson International Center for Scholars.
- Kontek, K. and Lewandowski, M. (2017). Range-Dependent Utility. *Management Science*, In Press.
- Kotler, P. and Zaltman, G. (1971). Social Marketing: An Approach to Planned Social Change. *The Journal of Marketing*, 35(3):3–12.
- Krastev, I. (2019). Putins Next Playground or the E.U.s Last Moral Stand? <https://www.nytimes.com/2019/01/28/opinion/russia-eu-balkans.html>. Retrieved on 29 January 2018.
- Kugler, T., Connolly, T., and Ordóñez, L. D. (2012). Emotion, Decision, and Risk: Betting on Gambles versus Betting on People. *Journal of Behavioral Decision Making*, 25(2):123–134.
- Kuhn, G. A. (2010). Roots of Peace Crop Income Projection: Afghanistan 2010. Roots of Peace.
- Kydd, A. H. (2007). *Trust and Mistrust in International Relations*. Princeton University Press, Princeton, NJ.
- Kydd, A. H. (2015). *International Relations Theory : The Game Theoretic Approach*. Cambridge University Press, Cambridge, UK.
- Larson, E. V., Darilek, R. E., Gibran, D., Nichiporuk, B., Richardson, A., Schwartz, L. H., and Thurston, C. Q. (2009). Foundations of Effective Influence Operations: A Framework for Enhancing Army Capabilities. Technical report, RAND Arroyo Center, Santa Monica, CA.
- Lee, D. R. and Clark, J. (2018). Can Behavioral Economists Improve Economic Rationality? *Public Choice*, 174(1-2):23–40.

- Leibowitz, M. L. and Lieberman, G. J. (1960). Optimal Composition and Deployment of a Heterogeneous Local Air-Defense System. *Operations Research*, 8(3):324–337.
- Lekivetz, R. and Jones, B. (2015). Fast Flexible Space-filling Designs for Nonrectangular Regions. *Quality and Reliability Engineering International*, 31(5):829–837.
- Lempert, K. M. and Phelps, E. A. (2014). Neuroeconomics of Emotion and Decision Making. In Glimcher, P. W. and Fehr, E., editors, *Neuroeconomics: Decision Making and the Brain*, chapter 12, pages 219–236. Elsevier, Oxford, UK.
- Lempert, R. J., Bryant, B. P., and Bankes, S. C. (2008). Comparing Algorithms for Scenario Discovery. Technical report, Santa Monica, CA: RAND Corporation.
- Lempert, R. J. and Collins, M. T. (2007). Managing the Risk of Uncertain Threshold Responses: Comparison of Robust, Optimum, and Precautionary Approaches. *Risk Analysis*, 27(4):1009–1026.
- Lempert, R. J., Groves, D. G., Popper, S. W., and Bankes, S. C. (2006). A General, Analytic Method for Generating Robust Strategies and Narrative Scenarios. *Management Science*, 52(4):514–528.
- Leppänen, I., Hämäläinen, R. P., Saarinen, E., and Viinikainen, M. (2018). Intrapersonal Emotional Responses to the Inquiry and Advocacy Modes of Interaction: a Psychophysiological Study. *Group Decision and Negotiation*, pages 1–16.
- Lerner, J. S. and Keltner, D. (2000). Beyond Valence: Toward a Model of Emotion-specific Influences on Judgement and Choice. *Cognition & Emotion*, 14(4):473–493.
- Lerner, J. S., Li, Y., Valdesolo, P., and Kassam, K. S. (2015). Emotion and Decision Making. *Annual Review of Psychology*, 66.
- Lerner, J. S., Small, D. A., and Loewenstein, G. (2004). Heart Strings and Purse Strings: Carryover Effects of Emotions on Economic Decisions. *Psychological Science*, 15(5):337–341.
- LeVeck, B. L., Hughes, D. A., Fowler, J. H., Hafner-Burton, E., and Victor, D. G. (2014). The Role of Self-interest in Elite Bargaining. In *Proceedings of the National Academy of Sciences*, volume 111, pages 18536–18541.
- Levy, J. S. (2003). Applications of Prospect Theory to Political Science. *Synthese*, 135(2):215–241.
- Lewis, A. S. and Overton, M. L. (2013). Nonsmooth Optimization via Quasi-Newton Methods. *Mathematical Programming*, 141(1-2):135–163.
- Lewis, R. M. and Torczon, V. (1999). Pattern Search Algorithms for Bound Constrained Minimization. *SIAM Journal on Optimization*, 9(4):1082–1099.

- Little, A. T. and Zeitzoff, T. (2017). A Bargaining Theory of Conflict with Evolutionary Preferences. *International Organization*, 71(3):523–557.
- Liu, Y.-J., Tsai, C.-L., Wang, M.-C., and Zhu, N. (2010). Prior Consequences and Subsequent Risk Taking: New Field Evidence from the Taiwan Futures Exchange. *Management Science*, 56(4):606–620.
- Louie, K. and De Martino, B. (2014). The Neurobiology of Context-Dependent Valuation and Choice. In Glimcher, P. W. and Fehr, E., editors, *Neuroeconomics: Decision Making and the Brain*, chapter 24, pages 455–476. Elsevier, San Diego, CA.
- Mansfield, D. (2017). Truly Unprecedented: How the Helmand Food Zone Supported an Increase in the Province’s Capacity to Produce Opium. Afghanistan Research and Evaluation Unit.
- Mansfield, D. (2018). Bombing Heroin Labs in Afghanistan: The Latest Act in the Theatre of Counternarcotics. LSE International Drug Policy Unit.
- Mansfield, D. and Fishstein, P. (2016). Time to Move on: Developing an Informed Development Response to Opium Poppy Cultivation in Afghanistan. Afghanistan Research and Evaluation Unit.
- Martin, E. and Symansky, S. (2006). Macroeconomic Impact of the Drug Economy and Counter-Narcotics Efforts. In Buddenberg, Doris and Byrd, William A., editor, *Afghanistan’s Drug Industry: Structure, Functioning Dynamics, and Implications for Counter-Narcotics Policy*, chapter 2, pages 25–51. The World Bank.
- McCauley, C. and Moskalenko, S. (2008). Mechanisms of Political Radicalization: Pathways Toward Terrorism. *Terrorism and Political Violence*, 20(3):415–433.
- McClintock, B. (2018). Russian Information Warfare: A Reality That Needs a Response. <https://www.rand.org/blog/2017/07/russian-information-warfare-a-reality-that-needs-a.html>. Retrieved on 18 Apr 2018.
- McDermott, R. (2007). *Presidential Leadership, Illness, and Decision Making*. Cambridge University Press, New York, NY.
- McEneaney, W. M., Fitzpartick, B. G., and Lauko, I. G. (2004). Stochastic Game Approach to Air Operations. *IEEE Transactions on Aerospace and Electronic Systems*, 40(4):1191–1216.
- McKelvey, R. D. and Palfrey, T. R. (1995). Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior*, 10(1):6 – 38.
- McKelvey, R. D. and Palfrey, T. R. (1998). Quantal Response Equilibria for Extensive Form Games. *Experimental Economics*, 11(1):9 – 41.

- Mejia, D. and Restrepo, P. (2016). The Economics of the War on Illegal Drug Production and Trafficking. *Journal of Economic Behavior & Organization*, 126:255–275.
- Milgrom, P. R. (1981). Good News and Bad News: Representation Theorems and Applications. *The Bell Journal of Economics*, pages 380–391.
- Moon, C. and Souva, M. (2016). Audience Costs, Information, and Credible Commitment Problems. *Journal of Conflict Resolution*, 60(3):434–458.
- Moravčík, Matej and Schmid, Martin and Burch, Neil and Lisý, Viliam and Morrill, Dustin and Bard, Nolan and Davis, Trevor and Waugh, Kevin and Johanson, Michael and Bowling, Michael (2017). Deepstack: Expert-level Artificial Intelligence in Heads-up No-limits Poker. *Science*, 356(6337):508–513.
- Moreno-Sanchez, R., Kraybill, D. S., Thompson, S. R., et al. (2002). An Economic Analysis of Coca Eradication Policy in Colombia. In *2002 Annual Meeting, July 28-31, Long Beach, CA*, number 19833. American Agricultural Economics Association.
- Nguyen, T. H., Sinha, A., and Tambe, M. (2016a). Addressing Behavioral Uncertainty in Security Games: An Efficient Robust Strategic Solution for Defender Patrols. In *IEEE International Parallel and Distributed Processing Symposium*, pages 1831–1838.
- Nguyen, T. H., Sinha, A., and Tambe, M. (2016b). Conquering Adversary Behavioral Uncertainty in Security Games: An Efficient Modeling Robust Based Algorithm. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 4242–4243.
- Nicolle, H. (2010). Review of Country Strategy (Badakhshan and Takhar). *Afghanistan Opium Survey 2018: Cultivation and Production*. Samuel Hall.
- Onge, J. and Floresco, S. (2009). Dopaminergic Modulation of Risk-based Decision Making. *Neuropsychopharmacology*, 34(3):681.
- Pachur, T., Suter, R. S., and Hertwig, R. (2017). How the Twain Can Meet: Prospect Theory and Models of Heuristics in Risky Choice. *Cognitive Psychology*, 93:44–73.
- Pain, A. (2010). Afghanistan Livelihood Trajectories: Evidence from Badakhshan. Afghanistan Research and Evaluation Unit.
- Pamp, O., Dendorfer, F., and Thurner, P. W. (2018). Arm Your Friends and Save on Defense? The Impact of Arms Exports on Military Expenditures. *Public Choice*, 177(1-2):165–187.
- Park, C. W. and Lessig, V. P. (1981). Familiarity and its Impact on Consumer Decision Biases and Heuristics. *Journal of Consumer Research*, 8(2):223–230.

- Phillips, P. J. and Pohl, G. (2017). Terrorist Choice: A Stochastic Dominance and Prospect Theory Analysis. *Defence and Peace Economics*, 28(2):150–164.
- Pittel, K. and Rubbelke, D. T. (2012). Decision Processes of a Suicide Bomber - The Economics and Psychology of Attacking and Defecting. *Defence and Peace Economics*, 23(2):251–272.
- Plott, C. R. and Zeiler, K. (2005). The Willingness to Pay-Willingness to Accept Gap, the "Endowment Effect," Subject Misconceptions, and Experimental Procedures for Eliciting Valuations. *The American Economic Review*, 95(3):530–545.
- Polyakova, A. and Boyer, S. P. (2018). The Future of Political Warfare: Russia, the West, and the Coming Age of Global Digital Competition. *Foreign Policy at Brookings*.
- Powell, R. (1990). *Nuclear Deterrence Theory: The Search for Credibility*. Princeton University Press, Princeton, NJ.
- Powell, R. (2015). Nuclear Brinkmanship, Limited War, and Military Power. *International Organization*, 69(3):589–626.
- Prelec, D. (1998). The Probability Weighting Function. *Econometrica*, 66(3):497–527.
- Qian, P. Z. (2012). Sliced Latin Hypercube Designs. *Journal of the American Statistical Association*, 107(497):393–399.
- Qureshi, A. S. (2002). Water Resources Management in Afghanistan: The Issues and Options. International Water Management Institute.
- Rapoport, A. and Chammah, A. M. (1966). The Game of Chicken. *American Behavioral Scientist*, 10(3):10–28.
- Rathbun, B. C., Kertzer, J. D., Reifler, J., Goren, P., and Scotto, T. J. (2016). Taking Foreign Policy Personally: Personal Values and Foreign Policy Attitudes. *International Studies Quarterly*, 60(1):124–137.
- Roberson, B. (2006). The Colonel Blotto Game. *Economic Theory*, 29(1):1–24.
- Robinson, L., Helmus, T. C., Cohen, R. S., Nader, A., Radin, A., Magnuson, M., and Migacheva, K. (2018). Modern Political Warfare. https://www.rand.org/pubs/research_reports/RR1772.html. RAND Corporation, Santa Monica, CA.
- Rogers, B. W., Palfrey, T. R., and Camerer, C. F. (2009). Heterogeneous Quantal Response Equilibrium and Cognitive Hierarchies. *Journal of Economic Theory*, pages 1440–1467.
- Roth, A. E. and Erev, I. (1998). Modeling How People Play Games: Reinforcement Learning in Experimental Games with Unique Mixed Strategies. *American Economic Review*, 88(4):848–881.

- Roy, B. (2010). Robustness in Operational Research and Decision Aiding: A Multifaceted Issue. *European Journal of Operational Research*, 200(3):629–638.
- Rubio, M. and Van Hollen, C. (2018). DETER Act. <https://www.congress.gov/bill/115th-congress/senate-bill/2785>. Retrieved on 16 August 2018.
- Russell, B. (1959). *Common Sense and Nuclear Warfare*. Routledge Classics, New York, NY.
- Ryoo, H. S. and Sahinidis, N. V. (1995). Global Optimization of Nonconvex NLPs and MINLPs with Application in Process Design. *Journal of Global Optimization*, 19(5):551–566.
- Ryoo, H. S. and Sahinidis, N. V. (1996). A Branch-and-Reduce Approach to Global Optimization. *Journal of Global Optimization*, 8(2):107–138.
- Sahinidis, N. V. (2018). BARON User Manual v. 2018.12.26. <http://www.minlp.com/downloads/docs/baron%20manual.pdf>. The Optimization Firm, LLC.
- Saunders, J. J. (2001). *The History of the Mongol Conquests*. University of Pennsylvania Press, Philadelphia, PA.
- Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press, Cambridge, MA.
- Schelling, T. C. (1966). *Arms and Influence*. Yale University Press, New Haven CT.
- Schulreich, S., Heussen, Y. G., Gerhardt, H., Mohr, P. N., Binkofski, F. C., Koelsch, S., and Heekeren, H. R. (2014). Music-evoked Incidental Happiness Modulates Probability Weighting During Risky Lottery Choices. *Frontiers in Psychology*, 4(1):1–17.
- Selten, R. (1998). Features of Experimentally Observed Bounded Rationality. *European Economic Review*, 42:413–436.
- Shamir, E. (2010). The Long and Winding Road: The US Army Managerial Approach to Command and the Adoption of Mission Command (Auftragstaktik). *The Journal of Strategic Studies*, 33(5):645–672.
- SHAPE (2013). Comprehensive Operations Planning Directive. NATO Supreme Headquarters Allied Powers Europe.
- Shen, L. and Bigsby, E. (2013). The Effects of Message Features. In *The SAGE Handbook of Persuasion: Developments in Theory and Practice*, pages 20–35. SAGE Publications Sage CA: Los Angeles, CA.
- Shi, Y. and Lian, Z. (2016). Optimization and Strategic Behavior in a Passenger–Taxi Service System. *European Journal of Operational Research*, 249(3):1024–1032.

- Shiv, B., Loewenstein, G., Bechara, A., Damasio, H., and Damasio, A. R. (2005). Investment Behavior and the Negative Side of Emotion. *Psychological Science*, 16(6):435–439.
- Shoham, Y. and Leyton-Brown, K. (2008). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York, NY.
- SIGAR (2018a). Counternarcotics: Lessons from the U.S. Experience in Afghanistan. Special Inspector General for Afghanistan Reconstruction.
- SIGAR (2018b). Quarterly Report to the United States Congress. Special Inspector General for Afghanistan Reconstruction.
- Silverstein, J. (2017). North Korea and China Also Interfered in U.S. Election, Reince Priebus Says. <http://www.nydailynews.com/news/politics/north-korea-china-interfered-u-s-election-priebus-article-1.3312905>. Retrieved on 9 August 2018.
- Slantchev, B. L. (2005). Military Coercion in Interstate Crises. *American Political Science Review*, 99(4):533–547.
- Stahl, D. O. and Wilson, P. W. (1995). On Players' Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 10(1):218–254.
- Stewart, N., Chater, N., and Brown, G. D. (2006). Decision by Sampling. *Cognitive Psychology*, 53(1):1–26.
- Stewart, N., Reimers, S., and Harris, A. J. (2015). On the Origin of Utility, Weighting, and Discounting Functions: How They Get Their Shapes and How to Change Their Shapes. *Management Science*, 61(3):687–705.
- Stowell, J. (2018). What is Hybrid Warfare? <https://globalsecurityreview.com/hybrid-and-non-linear-warfare-systematically-erases-the-divide-between-war-peace/>. Retrieved on 29 January 2018.
- Tawarmalani, M. and Sahinidis, N. V. (2004). Global Optimization of Mixed-integer Nonlinear Programs: A Theoretical and Computational Study. *Mathematical Programming*, 99(3):563–591.
- Tawarmalani, M. and Sahinidis, N. V. (2005). A Polyhedral Branch-and-Cut Approach to Global Optimization. *Mathematical Programming*, 103(2):225–249.
- Taylor, J. G. (1978). Differential-game Examination of Optimal Time-sequential fire-support Strategies. *Naval Research Logistics*, 25(2):323–355.
- Taylor, S. E. (1982). The Availability Bias in Social Perception and Interaction. In Kahneman, D., Slovic, P., and Tversky, A., editors, *Judgment Under Uncertainty: Heuristics and Biases*, pages 190–200. Cambridge University Press.

- Thaler, R. (1985). Mental Accounting and Consumer Choice. *Marketing Science*, 4(3):199–214.
- Thaler, R. H. (1999). Mental Accounting Matters. *Journal of Behavioral Decision Making*, 12(3):183.
- Thaler, R. H. and Sunstein, C. R. (2009). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press, New Haven, CT.
- The Economist (2019). A Russian Propaganda Outlet Prospers in Turkey. <https://www.economist.com/europe/2019/03/02/a-russian-propaganda-outlet-prospers-in-turkey>. Retrieved on 28 Feb 2019.
- The Optimization Firm (2019). BARON Publications. <http://www.minlp.com/baron-publications>. Accessed 4 Jan 2019.
- Theoi Project (2017). Peitho. <http://www.theoi.com/Daimon/Peitho.html>. Retrieved on 15 Feb 2018.
- Thomas, M. D. (2018). Reapplying Behavioral Symmetry: Public Choice and Choice Architecture. *Public Choice*, pages 1–15.
- Tom, S. M., Fox, C. R., Trepel, C., and Poldrack, R. A. (2007). The Neural Basis of Loss Aversion in Decision-making under Risk. *Science*, 315(5811):515–518.
- Tomz, M., Weeks, J. L., and Yarhi-Milo, K. (2018). Public Opinion and Decisions about Military Force in Democracies. <https://globalpoverty.stanford.edu/sites/default/files/publications/WP1027.pdf>. Stanford University Working Paper.
- Tóth, M. and Chytilék, R. (2018). Fast, Frugal and Correct? An Experimental Study on the Influence of Time Scarcity and Quantity of Information on the Voter Decision Making Process. *Public Choice*, pages 1–20.
- Truth Initiative (2018). truth. <https://www.thetruth.com/>. Retrieved on 13 Feb 2018.
- Tuchman, B. W. (1962). *The Guns of August*. Random House Publishing, New York, NY.
- Tversky, A. and Fox, C. R. (1995). Weighing Risk and Uncertainty. *Psychological Review*, 102(2):269.
- Tversky, A. and Kahneman, D. (1981). The Framing of Decisions and the Psychology of Choice. *Science*, 211(4481):453–458.
- Tversky, A. and Kahneman, D. (1992). Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323.

- Tversky, A. and Koehler, D. J. (1994). Support Theory: A Nonextensional Representation of Subjective Probability. *Psychological Review*, 101(4):547.
- U.K. Ministry of Defense (2013). Joint Doctrine Publication 5-0: Campaign Planning.
- United States Government Accountability Office (2010). Hybrid Warfare. Technical Report GAO-10-1036R, United States House of Representatives.
- United States Joint Chiefs of Staff (2012). Joint Publication 3-13: Information Operations.
- United States Military Academy (2014a). *The West Point History of the Civil War*, volume 1. Simon and Schuster, New York, NY.
- United States Military Academy (2014b). *The West Point History of the Great War*, volume 1. Simon and Schuster, New York, NY.
- United States Military Academy (2015a). *The West Point History of the World War II*, volume 1. Simon and Schuster, New York, NY.
- United States Military Academy (2015b). *The West Point History of the World War II*, volume 2. Simon and Schuster, New York, NY.
- UNODC (2003). The Opium Economy in Afghanistan: An International Problem. United Nations Office on Drugs and Crime.
- UNODC (2018a). Afghanistan Opium Survey 2017: Challenges to Sustainable Development, Peace and Security. <https://www.unodc.org/documents/crop-monitoring/Opium-survey-peace-security-web.pdf>. United Nations Office on Drugs and Crime.
- UNODC (2018b). Afghanistan Opium Survey 2018: Cultivation and Production. *Afghanistan Opium Survey 2018: Cultivation and Production*. United Nations Office on Drugs and Crime.
- U.S. Department of Defense (2018). Summary of the 2018 National Defense Strategy of The United States of America. <https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf>. U.S. Department of Defense.
- U.S. Joint Chiefs of Staff (2017a). Joint publication 3-0: Joint operations.
- U.S. Joint Chiefs of Staff (2017b). Joint Publication 5-0: Joint Planning.
- U.S. Joint Chiefs of Staff (2018). Joint Doctrine Note 1-18: Strategy. https://www.jcs.mil/Portals/36/Documents/Doctrine/jdn_jg/jdn1_18.pdf?ver=2018-04-25-150439-540. U.S. Joint Chiefs of Staff.

- von Neumann, J. (1928). Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320.
- Wakker, P. P. (2010). *Prospect Theory: For Risk and Ambiguity*. Cambridge University Press, New York, NY.
- Ward, C. and Byrd, W. A. (2004). *Afghanistan's Opium Drug Economy*. World Bank Washington, DC.
- Ward, C., Mansfield, D., Oldham, P., and Byrd, W. A. (2008). Afghanistan: Economic Incentives and Development Initiatives to Reduce Opium Production. The World Bank.
- Weber, M. and Zuchel, H. (2005). How Do Prior Outcomes Affect Risk Attitude? Comparing Escalation of Commitment and the House-money Effect. *Decision Analysis*, 2(1):30–43.
- Weeks, J. L. (2008). Autocratic Audience Costs: Regime Type and Signaling Resolve. *International Organization*, 62(1):35–64.
- White, L., Burger, K., and Yearworth, M. (2016). Big Data and Behavior in Operational Research: Towards a Smart OR. In *Behavioural Operational Research: Theory, Methodology and Practice*, pages 177–193. Palgrave Macmillan, London, UK.
- Williams, B. G. (2008). Mullah Omar's Missiles: A Field Report on Suicide Bombers in Afghanistan. *Middle East Policy*, 15(4):26.
- Wittchen, H.-U. and Jacobi, F. (2011). The Size and Burden of Mental Disorders and Other Disorders of the Brain in Europe 2010. *European Neuropsychopharmacology*, 21(9):655–679.
- Woody, C. (2017). Heroin is Driving a Sinister Trend in Afghanistan. Heroin is driving a sinister trend in Afghanistan. Business Insider.
- Woody, C. (2018a). Afghanistan is Producing a Lot More Opium than Before the US Invasion. The US Just Can't Stop It. <https://www.businessinsider.com/the-us-cant-seem-to-cut-back-afghanistans-opium-production-2018-6>. Business Insider.
- Woody, C. (2018b). NATO's Top Officer Says We're Living with 'a More Blurred Line Between Peace and War' Thanks to New Russian Tactics. <https://www.businessinsider.com/nato-jens-stoltenberg-world-faces-blurred-line-between-peace-and-war-2018-9>. Retrieved on 29 January 2018.
- Zagare, F. C. and Kilgour, D. M. (2000). *Perfect Deterrence*, volume 72. Cambridge University Press, Cambridge, UK.

Zajonc, R. B. (2001). Mere Exposure: A Gateway to the Subliminal. *Current Directions in Psychological Science*, 10(6):224–228.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

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

1. REPORT DATE (DD-MM-YYYY) 13 June 2019		2. REPORT TYPE Doctoral Dissertation		3. DATES COVERED (From - To) Mar 2017 – Jun 2019	
4. TITLE AND SUBTITLE Persuasion, Political Warfare, and Deterrence: Behavioral and Behaviorally Robust Models				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Caballero, William N, Capt, USAF				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way Wright-Patterson AFB, OH 45433-7765				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT-ENS-DS-19-J-022	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) STRATCOM Attn: Jill Morrissett, J902 (jill.s.morrissett.civ@mail.mil) 4048 Higley Rd Dahlgren, VA 22448				10. SPONSOR/MONITOR'S ACRONYM(S) USSTRATCOM	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Distribution Statement A. Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This dissertation examines game theory models in the context of persuasion and competition wherein decisionmakers are not completely rational by considering two complementary threads of research. The first thread of research pertains to offensive and preemptively defensive behavioral models. Research in this thread makes three notable contributions. First, an offensive modeling framework is created to identify how an entity optimally influences a populace to take a desired course of action. Second, a defensive modeling framework is defined wherein a regulating entity takes action to bound the behavior of multiple adversaries simultaneously attempting to persuade a group of decisionmakers. Third, an offensive influence modeling framework under conditions of ambiguity is developed in accordance with historical information limitations, and we demonstrate how it can be used to select a robust course of action on a specific, data-driven use case. The second thread of research pertains to behavioral and behaviorally robust approaches to deterrence. Research in this thread makes two notable contributions. First, we demonstrate the alternative insights behavioral game theory generates for the analysis of classic deterrence games, and explicate the rich analysis generated from its combined use with standard equilibrium models. Second, we define behaviorally robust models for an agent to use in a normal form game under varying forms of uncertainty in order to inform deterrence policy decisions.					
15. SUBJECT TERMS Influence, Deterrence, Information Operations, Behavioral Economics, Bounded Rationality, Persuasion					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 251	19a. NAME OF RESPONSIBLE PERSON Dr. Brian J. Lunday, AFIT/ENS
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (include area code) (937) 255-6565; brian.lunday@afit.edu

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18