Robust Optimization, Game Theory, and Variational Inequalities

by

Michele Leslie Aghassi

Sc.B., Applied Mathematics, Brown University (1998)

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2005

Author ....
Sloan School of Management
August 11, 2005

Certified by ....
Dimitris Bertsimas
Boeing Professor of Operations Research
Sloan School of Management
Thesis Supervisor

Accepted by ....
James B. Orlin
Edward Pennell Brooks Professor of Operations Research
Sloan School of Management
Co-director, Operations Research Center

Robust Optimization, Game Theory, and Variational Inequalities

by

Michele Leslie Aghassi

Submitted to the Sloan School of Management
on August 11, 2005, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

We propose a robust optimization approach to analyzing three distinct classes of problems related to the notion of equilibrium: the nominal variational inequality (VI) problem over a polyhedron, the finite game under payoff uncertainty, and the network design problem under demand uncertainty.

In the first part of the thesis, we demonstrate that the nominal VI problem is in fact a special instance of a robust constraint. Using this insight and duality-based proof techniques from robust optimization, we reformulate the VI problem over a polyhedron as a single-level (and many-times continuously differentiable) optimization problem. This reformulation applies even if the associated cost function has an asymmetric Jacobian matrix. We give sufficient conditions for the convexity of this reformulation and thereby identify a class of VIs, of which monotone affine (and possibly asymmetric) VIs are a special case, which may be solved using widely-available and commercial-grade convex optimization software.

In the second part of the thesis, we propose a distribution-free model of incomplete-information games, in which the players use a robust optimization approach to contend with payoff uncertainty. Our "robust game" model relaxes the assumptions of Harsanyi's Bayesian game model, and provides an alternative, distribution-free equilibrium concept, for which, in contrast to *ex post* equilibria, existence is guaranteed. We show that computation of "robust-optimization equilibria" is analogous to that of Nash equilibria of complete-information games. Our results cover incomplete-information games either involving or not involving private information.

In the third part of the thesis, we consider uncertainty on the part of a mechanism designer. Specifically, we present a novel, robust optimization model of the network design problem (NDP) under demand uncertainty and congestion effects, and under either system-optimal or user-optimal routing. We propose a corresponding branch and bound algorithm which comprises the first constructive use of the price of anarchy concept. In addition, we characterize conditions under which the robust NDP reduces to a less computationally demanding problem, either a nominal counterpart or a single-level quadratic optimization problem. Finally, we present a novel traffic "paradox," illustrating counterintuitive behavior of changes in cost relative to changes in demand.

Thesis Supervisor: Dimitris Bertsimas
Title: Boeing Professor of Operations Research
Sloan School of Management

Acknowledgments

I am very grateful to a number of people for their support and encouragement in making this thesis possible and in making the last four years at MIT so enjoyable. In fact, it is difficult to bring the full extent of my appreciation to life in these written acknowledgments. Nonetheless, I will try.

First and foremost, I would like to thank my advisor, Dimitris Bertsimas, for his expert guidance and his seemingly limitless knowledge. Dimitris is one of the most respected and prolific scientists in operations research today, and I have been extremely fortunate, for this and many other reasons, to have had the chance to collaborate with him over the last four years.

From my first days at MIT, he challenged me to become increasingly self-reliant. This push toward independence, despite yielding some growing pains early on, ultimately allowed my abilities to fully blossom, fostered my maturity as a researcher, and provided me with a greater sense of satisfaction during the many productive periods in my work. I quickly came to sincerely enjoy the freedom Dimitris gave me to chart my own course in my research. Moreover, I greatly appreciated that he always came to the rescue, with creative and inspiring ideas, whenever I was stuck and needed help.

As a result of Dimitris' mentorship, and his contagious enthusiasm, optimism, and belief in himself, I have learned a number of important lessons, relevant not only to research but also to life in general. First, you should always trust your instincts and ideas; indeed, you are often most effective when you make decisions as if you have the final say. Second, even when the odds of an accomplishment seem somewhat slim, if you believe, in your heart, that it is possible, you should remain optimistic about your efforts; this positive attitude often itself contributes to and allows for success. Third, in order to make a significant contribution in an area, your knowledge of that area need not be infallible; there is no shame in making mistakes or in asking basic questions. Because of these insights, I have become a more independent and capable person.

In addition, I would like to thank Georgia Perakis, who has co-advised me on Chapters 3 and 5 of this thesis. An incredible communicator and teacher, as well as an accomplished scientist, Georgia has had an influence on me since my first year at MIT, when I completed her nonlinear optimization course. Moreover, through two seminar courses that she led, Georgia encouraged my exposure to cutting-edge developments in the robust optimization and price of anarchy literatures. Indeed, in the latter of these two seminar courses, I learned the basic theory of variational inequalities and thereby thought to connect this problem with ideas from robust optimization. Most importantly, I am grateful to Georgia for the invaluable expertise and insight she has brought to this thesis in the areas of variational inequalities, network equilibrium problems, the price of anarchy, and network design. As a result of her comments, suggestions, and feedback, the period during which she has co-advised me has been my most productive and most rewarding at MIT.

It has been a great pleasure collaborating with Dimitris and Georgia. They are not only dynamic researchers, but also very caring friends. I will miss our weekly meetings together.

I am also grateful to Rob Freund for the time and effort he has contributed, as a member of my thesis committee, to this work. The care and detail with which Rob reviewed my thesis, and the time that he took to discuss his comments with me are a reflection of his exemplary commitment to research and teaching. His recommendations have, no doubt, significantly strengthened this thesis. In addition, in Georgia's seminar class on robust optimization, he often gave impromptu mini-lectures, from which I gained valuable insight into the theory of duality. I am furthermore thankful to Rob for the caring, personal interest he took in my career, by sharing with me his own career-related experiences and wisdom.

Having thanked my thesis committee members, I also wish to thank the National Science Foundation Graduate Research Fellowship Program and the MIT Presidential Fellowship Program for their generous financial support.

In addition, I am grateful to Melvyn Sim and Dessi Pachamanova, alumni of the Operations Research Center (ORC) at MIT, for helping to foster my interest in robust optimization. Their joint work with Dimitris Bertsimas in robust optimization, particularly that in the context of polyhedral uncertainty sets, has been a great inspiration to me.

On a more personal note, the ORC community's familial warmth and caring have made my time at MIT more enjoyable than I could have ever predicted. I have made some wonderful friends here. I feel very lucky to have met Susan Martonosi and Carol Meyers, my two best friends in the program. We have been each other's cheering squad over the years and have also shared a lot of laughs and fun times together. Along with Susan and Carol, I entered the ORC with a class of very kind and social people — Élodie Adida, Agustin Bompadre, Lincoln Chandler, Victor Martínez-de-Albéniz, Lauren McCann, and Maxime Phomma. I have also appreciated the friendship of John Bossert, Laura Kang, Adam Mersereau (my squash partner!), Sanne De Boer, Dave Czerwinski (my other squash partner!), Carine Simon, Ping Xu, and Brian Youngblood (an honorary "ORCer"). I will especially treasure our fun memories from the Courtside karaoke bar. In addition, I would like to thank Dave Craft, Jeff Hawkins, Romy Shioda, Anshul Sood, and Joe Wu for their camaraderie and for showing me the ropes at the ORC. Moreover, I am grateful to Paulette Mosley, Laura Rose, Veronica Mignott, and Andrew Carvalho for making the ORC such a great place.

Outside MIT, I have a wonderful circle of friends who have enriched my life. I treasure my friendship with Tracy Pesin, whom I consider family. Not only did Tracy proofread my grad school applications, but she has always been ready with advice and support when I've needed it. With her pep talks, she could transform any bad mood or sense of discouragement into a big smile and a desire to skip through the streets throwing flowers. I would also like to thank Kevin Dahl, Kaori Ogino, Jess Leader, Val Tirella, Tim Farrell, Ron Dor, Steve Aucello, Julie Houghton, Erica LeBow, and Pamela and Belden Daniels for their friendship over the years. Moreover, I am fortunate to have had the company of Dietrich, Archie, and Henry Springer while I was working on this thesis.

Of course, without the love and support of my family, I could not have accomplished what I have and the achievement of this thesis would not be as sweet. My mother and father have been a constant source of unconditional love, and have, all my life, encouraged me and believed in me with great enthusiasm. They have also instilled in me a sincere desire to learn and a deep respect of academic work. Likewise, I am grateful to my brother David and my sister Vivi for their love and support.

Finally, I would like to thank my wonderful girlfriend Kim, who has brightened my world immeasurably with her love, her caring, her enthusiasm for the little things in life, her amazing sense of humor, and her keen intelligence. Meeting Kim was, without a doubt, one of the most fortunate outcomes of my time at MIT. I have appreciated her strength and steadiness in encouraging me during the difficult periods, and her joy and excitement during the celebratory ones. More than anyone outside my thesis committee, Kim has shown a deep and sincere curiosity about my work that has meant a lot to me. I hope that I enrich her life as much as she enriches mine.

Contents

1 Introduction **17**

1.1 Dealing with Data Uncertainty in Optimization 17

1.2 Contributions and Structure of the Thesis 19

1.2.1 Nominal Variational Inequalities over Polyhedra 19

1.2.2 Games under Payoff Uncertainty 20

1.2.3 Network Design under Demand Uncertainty 20

1.3 Notation 21

2 Review of Robust Linear Optimization **23**

3 Variational Inequalities through Convex Programming: Insights from Robust Optimization **27**

3.1 Introduction 27

3.1.1 Variational Inequalities and Solution Methods 27

3.1.2 Complexity of VI Algorithms and VI Solver Availability 30

3.1.3 Contributions and Structure of this Chapter 32

3.2 Applying Duality-Based Techniques to VIs 34

3.2.1 Reformulation of a VI as a Nominal, Single-level Optimization Problem 35

3.2.2 Convexity of Reformulations 40

3.2.3 Reformulation of an MPEC as a Single-Level Optimization Problem . 42

3.2.4 Examples Admitting Reformulation as Convex Programs 45

3.2.5 Connection with the Classical Gap Function 46

3.3 Concise Reformulations . . . 47

4 Robust Game Theory **51**

4.1 Introduction . . . 51

4.1.1 Finite Games with Complete Information . . . 51

4.1.2 Finite Games with Incomplete Information . . . 52

4.1.3 Contributions and Structure of this Chapter . . . 54

4.1.4 Notation . . . 56

4.2 A Robust Approach to Payoff Uncertainty in Games . . . 56

4.2.1 Precedents for a Worst-Case Approach . . . 57

4.2.2 Formalization of the Robust Game Model . . . 59

4.2.3 Why Combine Equilibrium and Worst-Case Notions? . . . 65

4.2.4 Interpretation of Mixed Strategies . . . 66

4.2.5 Examples of Robust Finite Games . . . 68

4.2.6 Nonexistence of *Ex Post* Equilibria . . . 70

4.3 Existence of Equilibria in Robust Finite Games . . . 71

4.4 Computing Sample Equilibria of Robust Games . . . 75

4.4.1 Review for Complete-Information, Finite Games . . . 76

4.4.2 Robust Finite Games . . . 77

4.5 Comparison of Robust and Bayesian Finite Games . . . 90

4.5.1 Equilibria Sets Are Generally Not Equivalent . . . 91

4.5.2 Sizes of Sets of Equilibria . . . 92

4.5.3 Zero-sum Becomes Non-fixed-sum under Uncertainty . . . 94

4.5.4 Symmetric Robust Games and Symmetric Equilibria . . . 94

4.6 Robust Games with Private Information . . . 97

4.6.1 Extension of Model . . . 97

4.6.2 Existence of Equilibria . . . 100

4.6.3 Computation of Equilibria . . . 105

4.7 Conclusions . . . 106

5 Robust Transportation Network Design in User- and System-Equilibrium Environments 109

5.1 Introduction . . . 109

5.1.1 The Network Equilibrium Problem . . . 111

5.1.2 The Price of Anarchy . . . 113

5.1.3 Network Design under Selfish Routing and Demand Uncertainty . . . 114

5.1.4 Contributions and Structure of this Chapter . . . 116

5.1.5 Notation . . . 118

5.2 Formulation of the Robust Network Design Problem . . . 118

5.2.1 Review of NDP Definitions and Classical NEP Results . . . 118

5.2.2 Modeling the NDP under Demand Uncertainty . . . 125

5.3 Properties of the Robust NDP . . . 128

5.3.1 Convexity . . . 128

5.3.2 Equilibrium Costs May Decrease When Demands Increase . . . 133

5.3.3 A Single-Level Optimization Reformulation of $\tau_{UO}(D; \mathbf{y})$ . . . 142

5.3.4 Difficulty of the Robust vs. the Nominal NDP . . . 147

5.4 An Approximate Solution of $RNDP_{UO}(D)$ Based on the Price of Anarchy . 148

5.4.1 Price of Anarchy Review . . . 149

5.4.2 An Approximation Result . . . 151

5.5 A Branch and Bound Algorithm for the Robust NDP . . . 155

5.5.1 Notation and Definitions . . . 156

5.5.2 Bounding the Performance of a Partial Solution's Descendants . . . 158

5.5.3 The Branch and Bound Algorithm . . . 160

5.6 A Single-level QCLP Reformulation of $RNDP_{UO}(D)$ . . . 169

5.6.1 Solving the NEP in Closed Form under Parallel Arcs . . . 170

5.6.2 Heavy Traffic Conditions . . . 176

5.6.3 A QCLP Reformulation of $RNDP_{UO}(D)$ . . . 178

5.7 Conclusions . . . 182

6 Conclusions **183**

A An Alternate Derivation of the VI Reformulation **187**

B Comparison of Notions of Network Equilibrium **189**

List of Figures

5-1 The set of tuples of demand vectors and corresponding SO or UO flows may be nonconvex. 129

5-2 $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ may be a nonconvex function of $\mathbf{d}$ over D. 132

5-3 Without monotonicity, $\zeta_{UO}(\mathbf{y}, d)$ may decrease when d increases. 141

5-4 A branch and bound subtree rooted at $\boldsymbol{\pi}$ 158

5-5 A simple example of $RNDP_{UO}(D)$ 164

5-6 Features of the exact branch and bound algorithm for $RNDP_{SO}(D)$ 166

5-7 Features of the heuristic branch and bound algorithm for $RNDP_{UO}(D)$ . . . 167

5-8 An NDP with parallel paths and separable and affine path cost functions . . 171

List of Tables

4.1 Sizes of multilinear systems for equilibria 83
4.2 Numerical results for instances of robust network routing game 88

Chapter 1

Introduction

1.1 Dealing with Data Uncertainty in Optimization

Traditionally, the field of optimization has considered problems with known and deterministic objective value and constraint data. Research has focused on classifying problem complexity and determining exact algorithms (for polynomially solvable problems) or approximation algorithms and heuristics (for NP-hard problems). In practical applications, a decision maker must often commit to a course of action in the face of either constraint or objective data uncertainty.

One crude approach to dealing with such uncertainty is to ignore it and to simply fix problem data at "nominal" values suspected to be "representative." Of course, if the realized data do not match those used in the model, the model's optimal solution may, in practice, be infeasible or drastically sub-optimal, depending on the nature of the uncertainty. An alternative approach is to fix the unpredictable values in the model and then apply sensitivity analysis. Unfortunately, this method addresses only data perturbations that are infinitesimal or sufficiently small.

In contrast, stochastic programming incorporates data uncertainty directly into the model. This approach assumes that the decision maker knows, *a priori*, a set of possibly realizable scenarios (i.e., values of the data) and the full distribution giving the likelihood of realization of each. It then suggests a solution which produces the lowest average cost, where the average is taken over the aforementioned distribution. More advanced stochastic programming methods may factor risk into the objective rather than considering only expectations.

Nevertheless, stochastic programming has several drawbacks. First, the truly risk-averse decision maker may desire a guaranteed bound on the cost and may find such a guarantee more attractive than a good average cost. Second, when the data distribution is sufficiently complex, stochastic programming methods fall victim to the "curse of dimensionality." An exact model may be computationally intractable and may even remain intractable under simplifying approximations, including discretization, that do not excessively sacrifice accuracy. Third, there are no rigorous, general results discussing realized feasibility and cost when the true distribution is different (even to a small extent) from the one supposed in the model. Fourth, and perhaps most importantly, in many real-world applications, full distributional information on the uncertain quantities may be unavailable. In some situations, including those in which the uncertainty is not fundamentally stochastic in nature, the very existence of such a distribution may make little sense.

Fortunately, in light of these difficulties, robust optimization takes an approach to data uncertainty much different than those taken by sensitivity analysis and stochastic optimization. In robust optimization, data uncertainty is modeled as deterministic and is characterized by an uncertainty set, whose shape may be a box, an ellipsoid, a polyhedron, etc. In a robust optimization problem, one optimizes over the set of solutions feasible under all realizations of data from the uncertainty set. In the case of objective data uncertainty, robust optimization selects the solution yielding the best cost under the corresponding least favorable data realization. In this way, a robust optimization problem is either a minimax or a maximin problem; it seeks solutions which are feasible and perform best under the "worst-case" conditions allowed for by the given uncertainty set.

Robust optimization is an attractive problem-solving tool for several reasons. Unlike sensitivity analysis, it incorporates data uncertainty directly into the model and can handle perturbations of any (bounded) size. Unlike stochastic optimization, robust optimization is geared toward providing feasibility guarantees and cost bounds. In addition, robust optimization often leads to computationally tractable problems and is not hobbled by the curse of dimensionality. In some cases, even if robust optimization model assumptions are not satisfied (e.g., the size of the allowable perturbation), one can still make high-probability statements about feasibility. Lastly, robust optimization does not require detailed distributional information about the uncertain quantities. For these reasons, it is a valuable approach to problems characterized by data uncertainty.

1.2 Contributions and Structure of the Thesis

In this thesis, we propose a robust optimization approach to analyzing three distinct classes of problems: the nominal variational inequality (VI) problem over a polyhedron, the finite game under payoff uncertainty, and the network design problem under demand uncertainty. While not all of these problems are characterized by data uncertainty, they share a connection with the concept of equilibrium. We begin the body of the thesis, in Chapter 2, with a brief review of key results from the literature on robust optimization.

1.2.1 Nominal Variational Inequalities over Polyhedra

In Chapter 3, we demonstrate that, although the nominal VI problem involves no data uncertainty, it is in fact a special instance of a robust constraint. Using this insight and the duality-based proof techniques from the robust optimization literature, we reformulate the variational inequality problem over a polyhedron as a single-level (and many-times continuously differentiable) optimization problem. This reformulation applies even if the associated cost function has an asymmetric Jacobian matrix. We give sufficient conditions for the convexity of this reformulation. Accordingly, we provide a new approach, and one that draws on results from convex optimization, to VI complexity analysis. More importantly, however, we thereby identify a class of VIs, of which monotone affine (and possibly asymmetric) VIs are a special case, which may be solved using widely-available and commercial-grade convex optimization software.

1.2.2 Games under Payoff Uncertainty

In Chapter 4, we turn our attention to a class of problems that do involve data uncertainty. In particular, we propose a distribution-free model of incomplete-information games, in which the players use a robust optimization approach to contend with payoff uncertainty. Our "robust game" model relaxes the assumptions of Harsanyi's Bayesian game model, and provides an alternative, distribution-free equilibrium concept, which we call the "robust-optimization equilibrium," to that of the *ex post* equilibrium. We prove that the robust-optimization equilibria of an incomplete-information game subsume the *ex post* equilibria of the game and are, unlike the latter, guaranteed to exist when the game is finite and has bounded payoff uncertainty set. For arbitrary robust finite games with bounded polyhedral payoff uncertainty sets, we show that one may solve for a robust-optimization equilibrium in much the same way one would solve for a Nash equilibrium of a finite game with complete information. Our results cover incomplete-information games without private information as well as those involving potentially private information.

1.2.3 Network Design under Demand Uncertainty

Whereas, in Chapter 4, we consider a predetermined game mechanism and model the players as using robust optimization to deal with payoff uncertainty, in Chapter 5, we consider a mechanism design problem in which the designer is uncertain of conditions affecting the mechanism's performance. Specifically, we suggest a novel approach, based on robust optimization, to the binary choice, arc construction network design problem (NDP), under demand uncertainty, congestion effects, and either system-optimal (SO) or user-optimal (UO) routing. We propose a branch and bound algorithm for solving the resulting robust NDP. This algorithm comprises the first constructive use of the price of anarchy concept, which has previously been employed only in a descriptive, rather than a prescriptive manner. Moreover, using the notion of the price of anarchy, we prove that the optimal solution of the robust NDP under SO routing is an approximately optimal solution to the robust NDP under UO routing. In addition, we present conditions under which the robust NDP reduces to a less computationally demanding problem, either a nominal counterpart or a single-level quadratic optimization problem. Furthermore, we observe counterintuitive behavior, not yet noted in the literature, of costs at equilibrium with respect to changes in traffic demands on the network. The examples we present are analogous to Braess' Paradox [27] and illustrate that an increase in traffic demands on a network may yield a strict decrease in the costs at equilibrium. Finally, we establish convexity and monotonicity properties of functions relating to the worst-case performance of a given network design decision.

1.3 Notation

We will use the following notation conventions throughout the thesis. Boldface letters will denote vectors and matrices. In general, unless otherwise specified, upper case letters will signify matrices, while lower case will denote vectors. We will use subscripts to denote elements of a vector or matrix and superscripts to denote one entire vector or matrix in a sequence. To designate uncertain coefficients and their nominal counterparts, we will use the tilde (e.g., $\tilde{a}$) and the check (e.g., $\check{a}$), respectively. For a square, but not necessarily symmetric matrix $\mathbf{A}$, $\mathbf{A} \succeq \mathbf{0}$ will denote that $\mathbf{A}$ is positive semidefinite. Likewise, $\mathbf{A} \succ \mathbf{0}$ will denote that $\mathbf{A}$ is positive definite.

Chapter 2

Review of Robust Linear Optimization

For the purpose of more precisely characterizing the robust optimization approach, let us consider the mathematical optimization problem (MP)

$$
\begin{aligned}
P: \ \min_{\mathbf{x}} \quad & f(\mathbf{x}) \\
\text{s.t.} \quad & \mathbf{x} \in X(\check{\delta}_1, \check{\delta}_2, \ldots, \check{\delta}_\omega),
\end{aligned}
$$

where $\mathbf{x}$ is the vector of decision variables, $X(\delta_1, \delta_2, \ldots, \delta_\omega)$ is the feasible region defined by parameters δ_i, $i \in \{1, \ldots, \omega\}$, and $f(\mathbf{x})$ is the objective function. In the above nominal MP, we regard the parameter values as being fixed at $\delta_i = \check{\delta}_i$, $i \in \{1, \ldots, \omega\}$. Throughout the rest of this thesis, we use the term "nominal" to denote problems defined by fixed parameter values that are known with certainty.

Suppose instead that we do not know the exact values of these parameters $\tilde{\delta}_1, \ldots, \tilde{\delta}_\omega$, but know that $(\tilde{\delta}_1, \ldots, \tilde{\delta}_\omega)$ belongs to some uncertainty set U. Under this model of uncertainty, the robust counterpart of the above nominal problem is given by

$$\begin{array}{llll} RP: & \min\limits_{\mathbf{x}} & f(\mathbf{x}) & \\ & \text{s.t.} & \mathbf{x} \in X(\tilde{\delta}_1, \tilde{\delta}_2, \ldots, \tilde{\delta}_\omega), & \forall\, (\tilde{\delta}_1, \tilde{\delta}_2, \ldots, \tilde{\delta}_\omega) \in U. \end{array}$$

Without loss of generality, we can restrict our discussion of data uncertainty to the constraints, since the objective can always be incorporated into the constraints (see, e.g., Section 2 of [18]).

Initial results on robust linear optimization were given by Soyster in [150]. Soyster considered linear optimization problems (LPs) subject to column-wise uncertainty in the constraint matrix. His model is equivalent to the LP in which all uncertain parameters have been fixed at their worst-case values from the uncertainty set. Soyster's model is overly conservative; in practice, it seems quite unlikely that the uncertain parameters would all simultaneously realize their worst-case values. In addition, his model is specific to column-wise uncertainty and does not easily generalize.

Twenty years later, Ben-Tal and Nemirovski [11, 12, 13] and, independently, El Ghaoui et al. [45, 46], renewed the discussion of optimization under uncertainty. They examined ellipsoidal models of uncertainty, which, for the robust LP case, are less conservative than the column-wise model considered by Soyster. They showed that the robust counterpart of an LP under such ellipsoidal uncertainty models is a second-order cone optimization problem (SOCP). Furthermore, they remarked that polyhedral uncertainty can be regarded as a special case of ellipsoidal uncertainty. As a result, LPs under polyhedral uncertainty of the coefficient matrix can be solved via SOCPs.

Ellipsoidal uncertainty formulations of robustness are attractive in that they offer a reduced level of conservatism, as compared with the Soyster model, and lead to efficient solutions, via SOCPs, of LPs under uncertainty. Unfortunately, ellipsoidal uncertainty formulations give rise to robust counterparts whose exact solution is more computationally demanding than that of the corresponding nominal problem. In response to this drawback, Bertsimas and Sim [17, 18] offered an alternative model of uncertainty, under which the robust counterpart of an LP is an LP. Their formulation yields essentially the same level of conservatism as do those of Ben-Tal and Nemirovski and El Ghaoui et al.

Bertsimas, Pachamanova, and Sim [16] further extended the results of Bertsimas and Sim [18] to the case of general polyhedral uncertainty of the coefficient matrix. In particular, where $\tilde{\mathbf{A}}$ is an $m \times n$ coefficient matrix, they modeled the uncertainty set as a polyhedron

$$U = \left\{ \mathrm{vec}(\tilde{\mathbf{A}}) \;\middle|\; \mathbf{G} \cdot \mathrm{vec}(\tilde{\mathbf{A}}) \leq \mathbf{d} \right\}, \tag{2.1}$$

where $G \in \mathbb{R}^{\ell \times (mn)}$, $\mathbf{d} \in \mathbb{R}^{\ell}$, and $\mathrm{vec}(\tilde{\mathbf{A}}) \in \mathbb{R}^{(mn) \times 1}$ denotes the column vector obtained by stacking the row vectors of the matrix $\tilde{\mathbf{A}}$ one on top of another. They considered the robust LP given by

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{c}'\mathbf{x} \\ \text{s.t.} \quad & \tilde{\mathbf{A}}\mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \in S \\ & \forall \tilde{\mathbf{A}} \in U, \end{aligned} \tag{2.2}$$

where $\mathbf{x} \in \mathbb{R}^n$ is the decision vector and S is a polyhedron defined by constraints that are not subject to uncertainty. They showed that, if the number of constraints defining S is r, then robust LP (2.2) in n variables and $m + r$ constraints is equivalent to a nominal LP in $n + m\ell$ variables and $m^2 n + m + m\ell + r$ constraints.

A key element in the approach of Bertsimas et al. [16, 17, 18] is the use of dual variables and duality to reformulate robust constraints as equivalent nominal constraints over a modestly higher-dimensional variable space.

Chapter 3

Variational Inequalities through Convex Programming: Insights from Robust Optimization

3.1 Introduction

3.1.1 Variational Inequalities and Solution Methods

The variational inequality (VI) problem has engaged members of not only the optimization community, but also the mathematics, transportation science, engineering, and economics communities. Stated formally, given a set $K \subseteq \mathbb{R}^n$ and a mapping $\mathbf{F} : K \to \mathbb{R}^n$, the VI problem, denoted $VI(K, \mathbf{F})$, is to find an $\mathbf{x}^* \in K$ such that

$$\mathbf{F}(\mathbf{x}^*)'(\mathbf{x} - \mathbf{x}^*) \geq 0, \quad \forall \mathbf{x} \in K.$$

Variational inequalities subsume many other well-studied mathematical problems, including the solution of systems of equations, complementarity problems, and a class of fixed point problems. In addition, for any optimization problem over a closed, convex feasible region, the first-order optimality conditions comprise a VI. Accordingly, the VI problem also generalizes convex optimization.

Stampacchia and his collaborators [67, 92, 103, 151, 152] first introduced VIs as a tool for analyzing partial differential equations arising in mechanics. These early contributions, including the seminal 1966 paper by Hartman and Stampacchia [67], focused on infinite-dimensional VIs. In contrast, the study of finite-dimensional VIs grew out of the optimization community's interest in nonlinear complementarity problems, first introduced by Cottle [32, 33] in 1964. Soon thereafter, Lemke [89] extended the Lemke-Howson algorithm [90] to the linear complementarity problem, and Scarf [145] gave a method for computing the fixed points of a continuous mapping. These computational strides further bolstered the optimization community's theoretical interest in complementarity problems. In 1972, Karamardian [80] noted that the finite-dimensional complementarity problem is a special case of the finite-dimensional VI problem.

Just as the VI problem is of interest to widely ranging research communities, its corresponding computational solution methods are diverse. For a more complete discussion and history of these computational approaches than the one we give in this section, we refer the interested reader to the recent survey text by Facchinei and Pang [47], the monograph by Patriksson [124], and the references in both texts. The survey article by Harker and Pang [64] and the Ph.D. thesis of Hammond [63], as well as the references therein, also provide insightful reviews of the VI problem and associated algorithms.

One class of techniques for solving the VI problem consists of equation-based methods. These methods exploit the fact that $\mathbf{x}^*$ solves $VI(K, \mathbf{F})$ iff it satisfies a system of nonsmooth equations. In one such approach, the system of equations formulates the VI problem's solution set as the set of fixed points of a projection operator (see, e.g., Proposition 1.5.8 of [47], which is originally due to Eaves [44]). In order to solve the VI problem, one may thus use Newton methods or globally convergent techniques for solving such systems of nonsmooth equations.

Alternatively, the KKT conditions of a VI comprise a mixed complementarity problem (MiCP), involving both equations and nonnegativity constraints. The MiCP is equivalent, through a "complementarity function," to a system of equations alone, to which the aforementioned Newton methods apply. Rather than transforming the MiCP representation of the KKT conditions into a system of equations, one may treat them as a system of "constrained equations" and use interior point methods and smoothing techniques to handle this potentially inequality-constrained system. These approaches are modified Newton methods, tailored so that iterates satisfy a given constraint set. Interior point methods for solving VIs are similar, although not identical, to the celebrated interior point algorithms for solving linear and nonlinear convex optimization problems.

In contrast, instead of reformulating the VI problem as a system of equations or constrained equations, other contributors to the VI literature have suggested reformulating the VI as an equivalent optimization problem. Most simply, it is well known (see, e.g., Theorem 1.3.1 of Facchinei and Pang [47]) that if the Jacobian of $\mathbf{F}$, denoted by $J\mathbf{F}(\mathbf{x})$, is symmetric $\forall \mathbf{x} \in K$, then there exists a function $f : K \rightarrow \mathbb{R}$ such that

$$\nabla f(\mathbf{x}) \;=\; \mathbf{F}(\mathbf{x}), \quad \forall \mathbf{x} \in K.$$

If, in addition, K is convex, then $\mathbf{x}^*$ solves $VI(K, \mathbf{F})$ iff it is a stationary point of

$$\min_{\mathbf{x} \in K} \; f(\mathbf{x}). \tag{3.1}$$

When $\mathbf{F}$ is furthermore monotone over K, or equivalently, when its Jacobian matrix is positive semidefinite over K, the resulting optimization problem (3.1) is convex, and $\mathbf{x}^*$ solves $VI(K, \mathbf{F})$ iff it is an optimal solution of (3.1). Depending on the exact form of this convex program, one may apply any of a number of commercial solvers, including ILOG CPLEX, ILOG Solver, Xpress-MP, and LINDO. Moreover, when $\mathbf{F}$ is affine and K is polyhedral, problem (3.1) is a linearly constrained quadratic program (LCQP).

For a VI with possibly asymmetric $J\mathbf{F}(\mathbf{x})$, the concept of merit functions provides for the reformulation as an equivalent optimization problem. For $X \supseteq K$, where X is a closed set, a merit function for $VI(K, \mathbf{F})$ is defined to be a nonnegative $\theta : X \rightarrow \mathbb{R}$ such that $\mathbf{x}^*$

solves $VI(K, \mathbf{F})$ iff $\mathbf{x}^* \in X$ and $\theta(\mathbf{x}^*) = 0$. Thus, $VI(K, \mathbf{F})$ is equivalent to $\min_{\mathbf{x} \in X} \theta(\mathbf{x})$. Of course, if $X = \mathbb{R}^n$, then this reformulation yields an unconstrained optimization problem. The earliest and most intuitive merit function, proposed by Zuhovickii, Poljak, and Primak [174, 175, 176, 177] in the context of game-theoretic equilibrium computation, is the classical primal gap function, given by

$$\theta_{\text{gap}}(\mathbf{x}) \triangleq \sup_{\mathbf{y} \in K} \mathbf{F}(\mathbf{x})' (\mathbf{x} - \mathbf{y}). \tag{3.2}$$

Despite its simplicity, the literature has deprecated the use of this gap function in practice, because it may not be differentiable. Until 1989, the VI community regarded as an open question the issue of whether there exists a continuously differentiable optimization problem that equivalently reformulates $VI(K, \mathbf{F})$. In that year Auchmuty [3], and independently and soon thereafter Fukushima [55], proposed the regularized gap function, which altered the classical gap function in a way that guaranteed continuous differentiability. Others have since proposed numerous and varied merit functions, e.g., Wu, Florian and Marcotte [170], Zhu and Marcotte [173], Taji and Fukushima [156], Yamashita and Fukushima [171], Peng [125], and Yamashita, Taji, and Fukushima [172]. These merit functions give rise to specialized iterative descent methods for their minimization and, as a consequence, for the solution of the VI problem.

3.1.2 Complexity of VI Algorithms and VI Solver Availability

While the favored reformulations and associated algorithms for the VI problem have proved to be practically useful, they are not without drawbacks. For instance, the merit function approach to solving VIs involves an iterative descent method in which, at each iteration, evaluation of the merit function and its gradient requires the computation of a projection onto K. Under certain conditions, usually involving continuity and some version of monotonicity, iterative descent methods for merit function minimization yield limit points that are stationary points of the corresponding optimization problem (see, e.g., Section 10.6 in

Facchinei and Pang [47] for an overview). However, even when these requirements are met, the iterative descent algorithms may not converge in finite, let alone polynomial time.

In fact, the topic of complexity is not as often discussed in the VI literature as in other areas of optimization. Most algorithms are guaranteed to yield asymptotic convergence, with the exceptions of ellipsoid and interior point methods for VIs. In 1985, Lüthi [97] offered a specialized ellipsoid algorithm for the VI problem and proved its polynomial-time convergence for strongly monotone VIs. Perakis [126] and Magnanti and Perakis [101] proposed a class of ellipsoid-type algorithms ensuring polynomial-time convergence for a larger class of VIs, namely those that are strongly-$\mathbf{F}$-monotone. While these ellipsoid algorithms provide important insights into the complexity of the VI problem, they are not applied in practice.

Specialized interior point methods provide a more computationally useful and still efficient alternative to these ellipsoid algorithms. Indeed, polynomial-time interior point algorithms are available for certain classes of VI problems. For example, in the early 1990's, Harker and Xiao [65] proposed such an algorithm for the monotone affine VI problem, and Tseng [157] proposed another algorithm for the more general monotone VI problem. More recently, Sun and Zhao [154, 155] and Wu [168, 169] contributed efficient interior point methods for the monotone variational inequality with polyhedral feasible region K and $\mathbf{F}$ satisfying a "scaled Lipschitz condition."

Although interior point methods can efficiently solve interesting classes of VI problems, commercial-grade software packages for these algorithms are not available to the same extent as are commercial products for solving general convex QPs and other convex optimization problems. Accordingly, to solve instances of the VI problem, one would ideally like to harness the power of these industrial-strength convex optimization software packages. To the best of our knowledge, in the absence of symmetry of $J\mathbf{F}(\mathbf{x})$, the reformulation of a VI as a single-level, convex optimization problem, which is many-times continuously differentiable, is not well understood.

As already noted, the VI literature includes reformulations for problem instances in which $\mathbf{F}$ has a symmetric Jacobian over K. While the merit function technique further provides an optimization-based approach for the asymmetric case, merit functions are usually at most once continuously differentiable (see, e.g., the discussion in Section 10.6 of [47] or Section 4.2 of [48]). More importantly, for almost all merit functions, evaluation of the function itself requires solution of an optimization problem, generally involving a maximization over K. Thus, the merit function approach requires the solution of a bilevel optimization problem, a form that commercial optimization solvers cannot handle.

Finally, let us return to the case of a VI problem satisfying the symmetry and monotonicity requirements ensuring the existence of an objective function f, as in (3.1), inducing an equivalent convex optimization problem. Even in this case, identification of such an objective function f requires the evaluation, in closed form, of the indefinite line integral

$$f(\mathbf{x}) \;=\; \oint \mathbf{F}(\mathbf{x})' d\mathbf{x} \;=\; \int_0^1 \mathbf{F}\left(\mathbf{x}^0 + t\left(\mathbf{x} - \mathbf{x}^0\right)\right)' \left(\mathbf{x} - \mathbf{x}^0\right) dt,$$

where $\mathbf{x}^0$ is any vector in K. This integration is easy if $\mathbf{F}$ is affine. However, if $\mathbf{F}$ is nonlinear, and especially if n, the dimension of the problem, is large, analytically evaluating such an integral may be impractical. Ideally, one would like to develop a reformulation of the VI problem that does not require such preprocessing in order to generate a convex optimization problem for input into a commercial solver.

3.1.3 Contributions and Structure of this Chapter

In this chapter of the thesis, we study instances of $VI(K, \mathbf{F})$ in which K is polyhedral. To give an overview, we provide an approach to reformulating such VIs as optimization problems. Our method avoids the difficulties discussed in Section 3.1.2 and is based on ideas from robust optimization.

In more detail, our contributions are as follows.

1. In Section 3.2.1, we identify that VIs are in fact special instances of robust constraints. We observe, in Section 3.2.3, that the same is true of mathematical programs with equilibrium constraints (MPECs).

2. In Section 3.2.1, we exploit this observation and utilize the duality-based proof techniques from robust optimization in order to reformulate the set of $|K|$ constraints (typically infinitely many) comprising the VI, as a system of finitely many constraints whose size is polynomial in n and m. $|K|$ denotes the cardinality of the set K, and m denotes the number of equalities in the standard-form representation of polyhedral set K. We use this system to show that, when K is polyhedral, one can always reformulate the VI, even under asymmetry of $J\mathbf{F}(\mathbf{x})$, as an equivalent, single-level optimization problem. The objective and constraint functions of this equivalent problem are, in contrast to most merit functions, continuously differentiable as many times as $\mathbf{F}(\mathbf{x})$ is. We extend this reformulation result to the general case in which K need not be a polyhedron but is defined in terms of an arbitrary cone, i.e., not necessarily the nonnegative orthant.

3. In Section 3.2.2, we give sufficient conditions for the convexity of the VI reformulation. These sufficient conditions provide a new approach, and one that draws on results from convex optimization, to VI complexity analysis. More importantly, however, these sufficient conditions clearly identify a class of VI problems to which commercial, single-level convex optimization solvers may be applied. Furthermore, specializing this result to the affine case, we demonstrate that our reformulation is always convex when $\mathbf{F}$ is both affine and monotone. Thus, a monotone affine VI problem can always be reformulated as a convex LCQP, even when $J\mathbf{F}(\mathbf{x})$ is asymmetric. To illustrate the sufficient conditions for convexity in the nonlinear case, in Section 3.2.4, we give examples of VIs that have asymmetric $J\mathbf{F}(\mathbf{x})$ and satisfy these conditions.

4. In Section 3.2.5, we demonstrate that our VI reformulation can be viewed as equivalent to the classical primal gap function (3.2). Thus, in order to obtain a merit-function-based formulation that is continuously differentiable, one need not regularize the gap function, as was previously thought.

5. In Section 3.3, we consider the fact that our reformulation technique results in a modest augmentation of the space of variables. We show that we may in fact reduce the number of variables in our formulation so that it is comparable to the original variable space and to the formulation (3.1).

3.2 Applying Duality-Based Techniques to VIs

Having reviewed, in Chapter 2, the robust optimization paradigm, it is clear that a VI constraint on $\mathbf{x}^*$, which is of the form

$$\mathbf{F}(\mathbf{x}^*)'(\mathbf{x} - \mathbf{x}^*) \geq 0, \quad \forall \mathbf{x} \in K, \tag{3.3}$$

is, in fact, a special case of a robust constraint, in which it is as if $\mathbf{x}$ is subject to uncertainty and known only to belong to K. That is, the robust optimization framework encompasses the VI problem.

Accordingly, we use the duality-based techniques from robust optimization (see, e.g., Bertsimas and Sim [18]), to transform the robust constraint that comprises the VI problem into a set of finitely many nominal constraints and, alternatively, into a corresponding optimization problem. We then give a sufficient condition for the convexity of this nominal constraint set and the optimization problem it induces.

As already noted in Section 3.1.3, we focus our attention in this chapter of the thesis on the class of VIs in which K is polyhedral. However, we also extend our results to more general settings. We make no assumptions on the existence of solutions to the VI problems we consider.

3.2.1 Reformulation of a VI as a Nominal, Single-level Optimization Problem

In the case of polyhedral K, the VI problem is equivalent to a robust linear constraint. Without loss of generality, we may restrict our attention to standard-form instances,

$$K = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{A}\mathbf{x} = \mathbf{b},\ \mathbf{x} \geq \mathbf{0}\} \neq \emptyset, \tag{3.4}$$

where $\mathbf{A} \in \mathbb{R}^{m\times n}$ and $\mathbf{b} \in \mathbb{R}^m$. Reformulating K in standard form may require parallel changes in $\mathbf{F}$. For instance, recasting nonpositive variables as nonnegative ones requires changes of sign in $\mathbf{F}$, and replacement of free variables by the difference of nonnegative variables requires augmenting the dimension of $\mathbf{F}$.

We may now state and prove a constraint equivalence.

Theorem 3.2.1. *Suppose that K is the nonempty polyhedron given by (3.4). Then, $\mathbf{x}^*$ solves $VI(K, \mathbf{F})$ iff $\exists \boldsymbol{\lambda}^* \in \mathbb{R}^m$ such that $(\mathbf{x}, \boldsymbol{\lambda}) = (\mathbf{x}^*, \boldsymbol{\lambda}^*)$ satisfies*

$$\begin{aligned} \mathbf{F}(\mathbf{x})'\mathbf{x} &= \mathbf{b}'\boldsymbol{\lambda} \\ \mathbf{A}\mathbf{x} &= \mathbf{b} \\ \mathbf{x} &\geq \mathbf{0} \\ \mathbf{A}'\boldsymbol{\lambda} &\leq \mathbf{F}(\mathbf{x}). \end{aligned} \tag{3.5}$$

Proof. As is well-known in the VI literature (see., e.g. Section 1.2 of [47]), by the definition of the VI problem, $\mathbf{x}^* \in K$ satisfies (3.3) iff the following relation holds.

$$\begin{aligned} \mathbf{F}(\mathbf{x}^*)'\mathbf{x}^* = \min_{\mathbf{x}} \quad & \mathbf{F}(\mathbf{x}^*)'\mathbf{x} \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} = \mathbf{b} \\ & \mathbf{x} \geq 0. \end{aligned} \tag{3.6}$$

That is, $\mathbf{x}^*$ must itself optimize the LP (3.6) it induces. This well-known observation is credited to Eaves [44], who originally noted this equivalence in the context of the complementarity problem.

In LP (3.6), in the spirit of robust optimization, $\mathbf{x}^*$ is treated as data and $\mathbf{x}$ is the vector of decision variables. Since $\mathbf{x}^*$ in this way parameterizes LP (3.6), we refer to this LP as $LP(\mathbf{x}^*)$. Its dual, to which we refer as $DLP(\mathbf{x}^*)$, is

$$\begin{aligned} \max_{\boldsymbol{\lambda}} \quad & \mathbf{b}'\boldsymbol{\lambda} \\ \text{s.t.} \quad & \mathbf{A}'\boldsymbol{\lambda} \leq \mathbf{F}(\mathbf{x}^*). \end{aligned} \tag{3.7}$$

Suppose that $\mathbf{x}^*$ solves $VI(K, \mathbf{F})$. Then $LP(\mathbf{x}^*)$ has bounded optimal value, given by $\mathbf{F}(\mathbf{x}^*)'\mathbf{x}^*$. By LP strong duality, its dual, $DLP(\mathbf{x}^*)$, is also feasible with bounded optimal value equal to that of $LP(\mathbf{x}^*)$. Let $\boldsymbol{\lambda}^*$ denote an optimal solution of $DLP(\mathbf{x}^*)$. Then $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ satisfies (3.5).

For the reverse direction, suppose that $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ satisfies system (3.5). Then, $\mathbf{x}^*$ and $\boldsymbol{\lambda}^*$ are primal and dual feasible for $LP(\mathbf{x}^*)$ and $DLP(\mathbf{x}^*)$, respectively. Since $\mathbf{F}(\mathbf{x}^*)'\mathbf{x}^* = \mathbf{b}'\boldsymbol{\lambda}^*$, by LP weak duality, $\mathbf{x}^*$ must be optimal for $LP(\mathbf{x}^*)$. Therefore, $\mathbf{x}^*$ solves $VI(K, \mathbf{F})$. □

Remark: Since the KKT conditions for an LP are closely related to its dual, it is possible to prove Theorem 3.2.1 using these KKT conditions, rather than the dual LP in (3.7). In Appendix A, we provide such an alternate proof.

Theorem 3.2.1 implies the following equivalence of the VI with an optimization problem.

Corollary 3.2.1. *Suppose that K is the nonempty polyhedron given by (3.4).* $\mathbf{x}^*$ *solves* $VI(K, \mathbf{F})$ *iff the following MP has optimal value zero and* $\exists \boldsymbol{\lambda}^* \in \mathbb{R}^m$ *such that* $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ *is an* *optimal solution.*

$$\begin{aligned} \min_{\mathbf{x},\boldsymbol{\lambda}} \quad & \mathbf{F}(\mathbf{x})'\mathbf{x} - \mathbf{b}'\boldsymbol{\lambda} \\ s.t. \quad & \mathbf{A}\mathbf{x} = \mathbf{b} \\ & \mathbf{x} \geq \mathbf{0} \\ & \mathbf{A}'\boldsymbol{\lambda} \leq \mathbf{F}(\mathbf{x}). \end{aligned} \tag{3.8}$$

Proof. To begin, note that $\forall\,(\mathbf{x},\boldsymbol{\lambda})$ feasible for MP (3.8), $\mathbf{x}$ is a feasible solution of $LP(\mathbf{x})$, and $\boldsymbol{\lambda}$ is a feasible solution of $DLP(\mathbf{x})$. In addition, the objective function of MP (3.8) represents the duality gap of this primal-dual pair of solutions. Consequently, if MP (3.8) is feasible, by LP weak duality, its objective value is always nonnegative. From this observation and the equivalence proved in Theorem (3.2.1), the result immediately follows. □

Remark: It is worth pausing to note that, if $VI(K,\mathbf{F})$ has no solution, then MP (3.8) either is infeasible or has a strictly positive optimal value. Consider any $\mathbf{x}^* \in K$ that does *not* solve $VI(K,\mathbf{F})$. While we make no assumptions on the boundedness of K, because $K \neq \emptyset$, there are only two possibilities for $LP\,(\mathbf{x}^*)$. If this LP is unbounded, then there does not exist a $\boldsymbol{\lambda} \in \mathbb{R}^m$ such that $(\mathbf{x}^*,\boldsymbol{\lambda})$ is feasible for MP (3.8). Otherwise, $LP\,(\mathbf{x}^*)$ must be bounded but have optimal value strictly less than $\mathbf{F}\,(\mathbf{x}^*)'\,\mathbf{x}^*$. In this case, $\forall\boldsymbol{\lambda} \in \mathbb{R}^m$ such that $(\mathbf{x}^*,\boldsymbol{\lambda})$ is feasible for MP (3.8), the corresponding objective value is, by weak duality, strictly positive.

Corollary 3.2.1 establishes that any VI with polyhedral K can *always* be reformulated as a nominal, single-level optimization problem, even under asymmetry of the Jacobian matrix of the associated $\mathbf{F}(\mathbf{x})$. Thus, any such problem is always equivalent to an optimization problem of a format compatible with commercial optimization software, which does not accept constraints that themselves involve optimization problems. Furthermore, the objective and constraint functions of our reformulation (3.8) are, in contrast to most merit functions in the VI literature, continuously differentiable as many times as $\mathbf{F}(\mathbf{x})$ is.

For the affine VI, Corollary 3.2.1 implies that this problem can always be reformulated as a linearly-constrained quadratic program (LCQP). This implication generalizes the well-known equivalence (see, e.g., Section 1.5.3 of [47] or [167]) between the linear complementarity problem and the quadratic program (QP).

We have thus far considered only polyhedral K. In fact, we may extend our VI reformulation result to the more general case, in which K is defined in terms of an arbitrary cone $C \subseteq \mathbb{R}^n$. In this setting, C may be a cone other than the nonnegative orthant, in terms of which all polyhedra are defined. Without loss of generality, consider

$$K \;=\; \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{A}\mathbf{x} = \mathbf{b},\ \mathbf{x} \succeq_C \mathbf{0}\} \;\neq\; \emptyset. \tag{3.9}$$

We use the convention that, for any cone C, $\mathbf{x} \succeq_C \mathbf{0}$ denotes that $\mathbf{x} \in C$, and $\mathbf{x} \succ_C \mathbf{0}$ denotes that $\mathbf{x} \in \text{int}(C)$, the interior of C. Let C^* denote the cone dual to C, i.e.,

$$C^* \;=\; \{\boldsymbol{\eta} \in \mathbb{R}^n \mid \boldsymbol{\eta}'\mathbf{x} \geq 0,\ \forall \mathbf{x} \in C\}.$$

We omit the proof of the following theorem, since it is analogous to that of Corollary 3.2.1. For a review of conic duality, we refer the interested reader to Ben-Tal and Nemirovski [14].

Theorem 3.2.2. *Consider $VI(K, \mathbf{F})$, where K is given by (3.9). Suppose the following MP has optimal value zero and $\exists \boldsymbol{\lambda}^* \in \mathbb{R}^m$ such that $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ is an optimal solution.*

$$\begin{aligned} \min_{\mathbf{x},\boldsymbol{\lambda}} \quad & \mathbf{F}(\mathbf{x})'\mathbf{x} - \mathbf{b}'\boldsymbol{\lambda} \\ s.t. \quad & \mathbf{A}\mathbf{x} \;=\; \mathbf{b} \\ & \mathbf{x} \;\succeq_C\; \mathbf{0} \\ & \mathbf{A}'\boldsymbol{\lambda} \;\preceq_{C^*}\; \mathbf{F}(\mathbf{x}). \end{aligned} \tag{3.10}$$

Then $\mathbf{x}^*$ solves $VI(K, \mathbf{F})$.

Conversely, suppose that $\mathbf{x}^*$ solves $VI(K, \mathbf{F})$, that $\exists \mathbf{x} \in K$ such that $\mathbf{x} \succ_C \mathbf{0}$, and that $\exists \boldsymbol{\lambda} \in \mathbb{R}^m$ such that $\mathbf{A}'\boldsymbol{\lambda} \prec_{C^*} \mathbf{F}(\mathbf{x}^*)$. Then, MP (3.10) has optimal value zero, and $\exists \boldsymbol{\lambda}^* \in \mathbb{R}^m$

such that $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ *is an optimal solution.*

Let $\mathbb{M}^{m,n} \subseteq \mathbb{R}^{mn}$ denote the space of $m \times n$ matrices, $\mathbb{S}^n \subseteq \mathbb{R}^{n^2}$ denote the space of symmetric $n \times n$ matrices, and $\mathbb{S}^n_+$ denote the self-dual cone of symmetric, positive semidefinite, $n \times n$ matrices. In addition, let us define the inner product of $\mathbf{X} \in \mathbb{M}^{m,n}$ with $\mathbf{Y} \in \mathbb{M}^{m,n}$ as

$$\mathbf{X} \bullet \mathbf{Y} \;=\; \sum_{i=1}^{m} \sum_{j=1}^{n} X_{ij} Y_{ij},$$

where X_{ij} and Y_{ij} are the $(i,j)^{\text{th}}$ elements of the matrices $\mathbf{X}$ and $\mathbf{Y}$, respectively. Setting $C = \mathbb{S}^n_+$ in Theorem 3.2.2, we obtain the following corollary, specific to VIs over subsets of $\mathbb{S}^n_+$. Note that, in terms of the notation conventions set forth in Section 1.3 and in this section, for $\mathbf{X} \in \mathbb{M}^{n,n}$, $\mathbf{X} \succeq_{\mathbb{S}^n_+}$ is a stronger statement than $\mathbf{X} \succeq \mathbf{0}$, since the former requires that $\mathbf{X}$ is symmetric, while the latter does not.

Corollary 3.2.2. *Consider an arbitrary* $\mathbf{F} : \mathbb{M}^{n,n} \to \mathbb{M}^{n,n}$ *and* $VI\,(K, \mathbf{F})$*, where*

$$K \;=\; \left\{ \mathbf{X} \in \mathbb{M}^{n,n} \;\middle|\; \mathbf{A}^i \bullet \mathbf{X} = b_i,\ i = 1, \ldots, m;\ \mathbf{X} \succeq_{\mathbb{S}^n_+} \mathbf{0} \right\}, \tag{3.11}$$

$\mathbf{A}^i \in \mathbb{M}^{n,n}$*, and* $b_i \in \mathbb{R}$*,* $i \in \{1, \ldots, m\}$*. Suppose the following MP has optimal value zero and* $\exists \boldsymbol{\lambda}^* \in \mathbb{R}^m$ *such that* $(\mathbf{X}^*, \boldsymbol{\lambda}^*)$ *is an optimal solution.*

$$\begin{aligned}
\min_{\substack{\mathbf{X} \in \mathbb{M}^{n,n} \\ \boldsymbol{\lambda} \in \mathbb{R}^m}} \quad & \mathbf{F}(\mathbf{X}) \bullet \mathbf{X} - \mathbf{b}'\boldsymbol{\lambda} & \\
s.t. \quad & \mathbf{A}^i \bullet \mathbf{X} \;=\; b_i, & i = 1, \ldots, m \\
& \mathbf{X} \;\succeq_{\mathbb{S}^n_+}\; \mathbf{0} & \\
& \sum_{i=1}^{m} \lambda_i \mathbf{A}^i \;\preceq_{\mathbb{S}^n_+}\; \mathbf{F}(\mathbf{X}). &
\end{aligned} \tag{3.12}$$

Then $\mathbf{X}^*$ *solves* $VI(K, \mathbf{F})$.

Conversely, suppose that $\mathbf{X}^*$ *solves* $VI(K, \mathbf{F})$*, that* $\exists \mathbf{X} \in K$ *such that* $\mathbf{X} \succ_{\mathbb{S}^n_+} \mathbf{0}$*, and that* $\exists \boldsymbol{\lambda} \in \mathbb{R}^m$ *such that* $\sum_{i=1}^m \lambda_i \mathbf{A}^i \prec_{\mathbb{S}^n_+} \mathbf{F}(\mathbf{X}^*)$*. Then, MP (3.12) has optimal value zero, and* $\exists \boldsymbol{\lambda}^* \in \mathbb{R}^m$ *such that* $(\mathbf{X}^*, \boldsymbol{\lambda}^*)$ *is an optimal solution.*

Remark: Note that when

$$\mathbf{F}(\mathbf{X}) = \sum_{i=1}^{n}\sum_{j=1}^{n} X_{ij}\mathbf{G}^{ij} + \mathbf{H},$$

where $\mathbf{G}^{ij} \in \mathbb{M}^{n,n}$, $\mathbf{H} \in \mathbb{M}^{n,n}$, $i, j \in \{1, \dots, n\}$, the objective function of MP (3.12) is quadratic in $\mathbf{X}$ and $\boldsymbol{\lambda}$.

3.2.2 Convexity of Reformulations

Having discussed a VI reformulation that uses duality-based techniques from robust optimization, we next use this reformulation to analyze the complexity of the original VI. In particular, converting the VI into a single-level optimization problem facilitates certification of polynomial-time complexity. Indeed, one need not bother proving the polynomial-time convergence of an algorithm specifically tailored to the VI. Rather, if our proposed optimization problem reformulation is solvable in polynomial time (e.g., if it is convex and satisfies other complexity conditions), then the complexity of the original VI is polynomial time. To understand the justification of this result, note that if the standard-form representation of a polyhedron K involves $m + n$ constraints and n variables, then VI reformulation (3.8) is defined by $m + 2n$ constraints in $m + n$ variables. In the following theorem, we give a sufficient condition for the convexity of this reformulation.

Theorem 3.2.3. *Suppose that K is the nonempty polyhedron given by (3.4). If $\mathbf{F}_j(\mathbf{x})$ is a concave function over K, $\forall j \in \{1, \dots, n\}$, and $\mathbf{F}(\mathbf{x})'\mathbf{x}$ is a convex function over K, then system (3.5) defines a convex feasible region and MP (3.8) defines a convex optimization problem.*

Proof. The result for MP (3.8) follows immediately from Corollary 3.2.1. If we replace the constraint $\mathbf{F}(\mathbf{x})'\mathbf{x} = \mathbf{b}'\boldsymbol{\lambda}$ in system (3.5) with $\mathbf{F}(\mathbf{x})'\mathbf{x} \leq \mathbf{b}'\boldsymbol{\lambda}$, we obtain a resulting system that is equivalent to (3.5). The reason is that, as explained in the proof of Corollary 3.2.1, $\mathbf{F}(\mathbf{x})'\mathbf{x} - \mathbf{b}'\boldsymbol{\lambda}$ is the duality gap of primal-dual pair (3.6) and (3.7) and therefore cannot be strictly negative. Accordingly, the convexity of the solution set of system (3.5) follows immediately from Theorem 3.2.1. □

Remark: It is worth noting that if K is not, without reformulation, a subset of the nonnegative orthant, then the concavity requirements on $F_j(\mathbf{x})$ in Theorem 3.2.3 may be directly stated as follows. If, $\forall \mathbf{x} \in K$, $x_j \geq 0$, then the sufficient condition asks that $F_j(\mathbf{x})$ be concave, as before. Alternatively, if, $\forall \mathbf{x} \in K$, $x_j \leq 0$, then the sufficient condition asks that $F_j(\mathbf{x})$ be convex. Otherwise, the sufficient condition asks that $F_j(\mathbf{x})$ be affine.

Returning to our discussion of the implications and significance of Theorem 3.2.3, if $VI(K, \mathbf{F})$ satisfies the theorem's conditions, it is equivalent to a convex optimization problem, given by MP (3.8), and can therefore be solved by commercial convex optimization software. For the affine case, the conditions of Theorem 3.2.3 simplify considerably. In particular, when $\mathbf{F}(\mathbf{x})$ is affine, the following equivalence holds. $\mathbf{F}_j(\mathbf{x})$ is a concave function over K, $\forall j \in \{1, \dots, n\}$, and $\mathbf{F}(\mathbf{x})'\mathbf{x}$ is a convex function over K iff $\mathbf{F}$ is monotone over K. The reason is that monotonicity holds iff $J\mathbf{F}(\mathbf{x}) \succeq \mathbf{0}$ (see, e.g., Proposition 2.3.2 of [47]). Accordingly, any affine monotone VI is polynomially solvable using a commercial QP solver applied to a convex LCQP. In contrast, when $\mathbf{F}(\mathbf{x})$ is not affine, the aforementioned equivalence fails. In particular, if $\mathbf{F}_j(\mathbf{x})$ is a concave function over K, $\forall j \in \{1, \dots, n\}$, but $\mathbf{F}(\mathbf{x})$ is not affine, then monotonicity on K is necessary, but not sufficient for $\mathbf{F}(\mathbf{x})'\mathbf{x}$ to be a convex function over K. For example, for $n = 1$, $K = [0, \infty)$, and $F(x) = -e^{-x} + 1$, $F(x)$ is concave and monotone over K, but $F(x) \cdot x$ is concave for $x \geq 2$.

We summarize the above discussion of the affine case in the following corollary.

Corollary 3.2.3. *Suppose that K is the nonempty polyhedron given by (3.4), and that $\mathbf{F}(\mathbf{x}) = \mathbf{G}\mathbf{x} + \mathbf{h}$, with $\mathbf{G} \succeq \mathbf{0}$, but not necessarily symmetric. Then, system (3.5) is a set of quadratic constraints defining a convex feasible region and MP (3.8) is a convex LCQP.*

Similarly, for the setting of $VI(K, \mathbf{F})$ with $K \subseteq \mathbb{S}^n_+$, we have the following analogous result.

Corollary 3.2.4. *Suppose that K is given by (3.11), and that*

$$\mathbf{F}(\mathbf{X}) = \sum_{i=1}^{n}\sum_{j=1}^{n} X_{ij}\mathbf{G}^{ij} + \mathbf{H},$$

where $\mathbf{G}^{ij} \in \mathbb{M}^{n,n}$, $\mathbf{H} \in \mathbb{M}^{n,n}$, $i,j \in \{1,\ldots,n\}$. Suppose that $\boldsymbol{\mathcal{G}} \succeq \mathbf{0}$, where

$$\boldsymbol{\mathcal{G}} = \left[vec\left(\mathbf{G}^{11}\right), vec\left(\mathbf{G}^{12}\right), \ldots, vec\left(\mathbf{G}^{1n}\right), \ldots, vec\left(\mathbf{G}^{n1}\right), \ldots, vec\left(\mathbf{G}^{nn}\right)\right],$$

and $vec\left(\mathbf{G}^{ij}\right)$ denotes the column vector obtained by stacking the row vectors of the matrix $\mathbf{G}^{ij}$ one on top of another. Then, MP (3.12) is convex, with a quadratic objective.

Proof. The result follows, since $\mathbf{F}(\mathbf{X}) = \text{vec}\,(\mathbf{X})'\,\boldsymbol{\mathcal{G}}\,\text{vec}\,(\mathbf{X}) + \text{vec}\,(\mathbf{H})$. □

Returning to the case of polyhedral K, in the next subsection, we extend the discussion to the mathematical program with equilibrium constraints (MPEC). Before doing so, let us pause to underscore, as we did in Section 3.1, that our contribution sheds additional light on VI complexity analysis by offering an insightful method — one that leverages the theory of convex optimization — of determining that a VI is efficiently solvable. However, we believe the appeal of our approach is due as much, if not more so, to the fact that it extends the applicability of commercial, single-level optimization solvers to a larger class of VIs. In contrast, the VI literature has focused on developing specialized iterative algorithms for these problems. As we noted in the introduction, commercial optimization packages for single-level problems are available, supported, and refined through practice to a greater extent than are solvers specific to VIs. Therefore, it is to one's benefit to use these commercial solvers if possible.

3.2.3 Reformulation of an MPEC as a Single-Level Optimization Problem

We now observe that the MPEC is, like the VI, a special case of the robust optimization problem. Let $SOL(K, \mathbf{F})$ denote the solution set of $VI(K, \mathbf{F})$. Using this notation, the

MPEC is the following optimization problem whose constraints include a parameterized VI.

$$\begin{aligned} \min_{\mathbf{u}\in\mathbb{R}^{n_1},\ \mathbf{x}\in\mathbb{R}^{n_2}} \quad & g(\mathbf{u},\mathbf{x}) \\ \text{s.t.} \quad & (\mathbf{u},\mathbf{x}) \ \in \ S \\ & \mathbf{x} \ \in \ SOL\,(K(\mathbf{u}),\mathbf{F}(\mathbf{u};\cdot))\,. \end{aligned} \tag{3.13}$$

$K(\mathbf{u})$ denotes a feasible region parameterized by the vector $\mathbf{u}$ of so-called upper-level decision variables, and $\mathbf{F}(\mathbf{u};\cdot)$ denotes a function, which is also parameterized by $\mathbf{u}$ and whose values are elements in $\mathbb{R}^{n_2}$. Analogously, $\mathbf{x}$ is called the vector of lower-level decision variables. Because the parameterized VI is a class of instances of robust constraints, it follows that the MPEC is a special case of the robust optimization problem. Accordingly, the VI reformulation results in Section 3.2.1 extend to MPECs.

In general, the exact value of $\mathbf{u}$ may determine not only the coefficients in the constraints defining $K(\mathbf{u})$ but also the number of constraints. Let us focus on instances in which $\mathbf{u}$ affects only the coefficients in the constraints, i.e., cases in which the following condition holds.

Condition 3.2.1. $\forall \mathbf{u} \in \mathbb{R}^{n_1}$ such that $\exists \mathbf{x} \in \mathbb{R}^{n_2}$ with $(\mathbf{u},\mathbf{x}) \in S$, $\exists m < \infty$ for which $K(\mathbf{u})$ is a nonempty polyhedron given in standard form by

$$K(\mathbf{u}) \ = \ \{\mathbf{x} \in \mathbb{R}^{n_2} \ | \ [\mathbf{A}(\mathbf{u})]\,\mathbf{x} = \mathbf{b}(\mathbf{u}),\ \mathbf{x} \geq \mathbf{0}\}\,,$$

with $[\mathbf{A}(\mathbf{u})] \in \mathbb{R}^{m\times n_2}$ and $\mathbf{b}(\mathbf{u}) \in \mathbb{R}^m$.

Theorem 3.2.1 yields the following corollary.

Corollary 3.2.5. *Consider MPEC (3.13) satisfying Condition 3.2.1. Then,* $(\mathbf{u}^*,\mathbf{x}^*)$ *is an optimal solution of MPEC (3.13) iff* $\exists \boldsymbol{\lambda}^* \in \mathbb{R}^m$ *such that* $(\mathbf{u},\mathbf{x},\boldsymbol{\lambda}) = (\mathbf{u}^*,\mathbf{x}^*,\boldsymbol{\lambda}^*)$ *is an optimal* *solution of the following problem.*

$$
\begin{aligned}
\min_{\mathbf{u},\mathbf{x},\boldsymbol{\lambda}} \quad & g(\mathbf{u},\mathbf{x}) \\
s.t. \quad & (\mathbf{u},\mathbf{x}) \in S \\
& \mathbf{F}(\mathbf{u};\mathbf{x})'\mathbf{x} = \mathbf{b}(\mathbf{u})'\boldsymbol{\lambda} \\
& [\mathbf{A}(\mathbf{u})]\,\mathbf{x} = \mathbf{b}(\mathbf{u}) \\
& \mathbf{x} \geq \mathbf{0} \\
& [\mathbf{A}(\mathbf{u})]'\,\boldsymbol{\lambda} \leq \mathbf{F}(\mathbf{u};\mathbf{x}).
\end{aligned}
\tag{3.14}
$$

The fact that the MPEC with polyhedral $K(\mathbf{u})$ may be converted into a single-level optimization problem is well known in the literature (see, e.g., [96]). The classical single-level MPEC reformulation is KKT-based, in that it dictates replacement of the VI constraint with its KKT conditions. In contrast, the derivation method used in Corollary 3.2.5 and in the proof of Theorem 3.2.1 is based on LP duality. Especially when $K(\mathbf{u})$ is not, without augmentation of the lower-level space of variables, in standard form, the duality-based reformulation may be appealing. In particular, although the two approaches are closely related, unlike its KKT-based analog, the duality-based reformulation does not give rise to complementarity constraints in the resulting single-level optimization equivalent of the MPEC. In some settings, this lack of complementarity constraints may facilitate identification of convexity of the MPEC reformulation and may provide greater flexibility for further manipulation of this reformulation.

In general, certifying the convexity of MPEC reformulation (3.14) is a bit more involved than doing the same for VI reformulation (3.5). Specifically, if $K(\mathbf{u})$ and $F(\mathbf{u};\cdot)$ truly depend on $\mathbf{u}$, in general, the reformulation will be nonconvex. For example, if the dependence on $\mathbf{u}$ is linear, then reformulation (3.14) will contain terms involving products of the upper- and lower-level decision variables. Otherwise, in the simple case in which $K(\mathbf{u})$ and $F(\mathbf{u};\cdot)$ are constant with respect to all feasible $\mathbf{u}$, reformulation (3.14) is convex if S is convex, if the upper-level objective g is convex, and if, $\forall \mathbf{u}$ feasible for the MPEC, $K(\mathbf{u})$ and $\mathbf{F}(\mathbf{u};\cdot)$ satisfy the conditions of Theorem 3.2.3.

3.2.4 Examples Admitting Reformulation as Convex Programs

Returning to the topic of VIs, in this subsection, we illustrate our VI reformulation technique with examples of $VI(K, \mathbf{F})$ satisfying the conditions of Theorem 3.2.3. Recall that these conditions are sufficient for the convexity of the VI reformulations given in Section 3.2.1. As noted in that section, if $\mathbf{F}(\mathbf{x})$ is affine, then $VI(K, \mathbf{F})$ is always equivalent to an LCQP, even if $J\mathbf{F}(\mathbf{x})$ is asymmetric. Moreover, this LCQP is convex as long as $\mathbf{F}(\mathbf{x})$ is monotone. Since it is obvious that there exist such instances involving monotone $\mathbf{F}(\mathbf{x})$ with asymmetric $J\mathbf{F}(\mathbf{x})$, we give examples only of instances with nonlinear $\mathbf{F}(\mathbf{x})$.

Example 1. Let $K \subseteq R^n_+$ be a bounded polyhedron, where $\mathbb{R}^n_+$ denotes the nonnegative orthant. Let $\mathbf{F}(\mathbf{x}) = \mathbf{H}(\mathbf{x}) + \mathbf{G}\mathbf{x} + \mathbf{h}$, where $\mathbf{G}$ is asymmetric, and

$$\begin{aligned} H_j(\mathbf{x}) &= \ln x_j, \qquad\qquad j = 1, \ldots, n \\ \mathbf{G} &\succeq \mathbf{0}. \end{aligned}$$

Because $\mathbf{F}$ is continuous and K is compact and convex, this VI problem is guaranteed to possess a solution (see, e.g., Corollary 2.2.5 of [47]).

Example 2. Consider the same setting as Example 1, but with

$$H_j(\mathbf{x}) = \sqrt{x_j}, \qquad\qquad j = 1, \ldots, n.$$

Example 3. Finally, consider any $K \subseteq \mathbb{R}^2_+ \backslash \{(0,0)\}$ such that $\left(\frac{1}{3}, \frac{1}{3}\right) \in K$. Let

$$\mathbf{F}(\mathbf{x}) = \begin{pmatrix} \ln(2x_1 + x_2) \\ \ln(x_1 + 2x_2) \end{pmatrix}.$$

Clearly, $\left(\frac{1}{3}, \frac{1}{3}\right)$ is a solution of this VI. By inspection, it is not obvious that this example satisfies the sufficient conditions of Theorem 3.2.3. However, one may symbolically compute, e.g., using MAPLE, the eigenvalues of the Hessians of $F_1(\mathbf{x})$, $F_2(\mathbf{x})$, and $\mathbf{F}(\mathbf{x})'\mathbf{x}$. In doing so, it becomes clear by inspection that, because $K \subseteq \mathbb{R}^2_+$, the signs of these eigenvalues over K imply the concavity of $F_j(\mathbf{x})$, $j \in \{1, 2\}$, and the convexity of $\mathbf{F}(\mathbf{x})'\mathbf{x}$ over K (see, e.g., Proposition B.4 of Bertsekas [15]).

3.2.5 Connection with the Classical Gap Function

Thus far, we have shown that the VI over a polyhedron can always be reformulated as a single-level optimization problem, even when the corresponding Jacobian is asymmetric. In addition, we gave sufficient conditions for the convexity of these equivalent optimization problems and thereby identified a class of VIs which can be solved using, as a black box, convex optimization and the widely available commercial solvers for such problems.

In this subsection, we consider a connection between our approach and the merit function approach, the latter of which induces bilevel programming reformulations of VIs. Recall the classical primal gap function,

$$\theta_{\text{gap}}(\mathbf{x}) \triangleq \sup_{\mathbf{y} \in K} \mathbf{F}(\mathbf{x})' (\mathbf{x} - \mathbf{y}) .$$

Interestingly, for K given by (3.4), our reformulation of $VI(K, \mathbf{F})$ as equivalent optimization problem (3.8), is in fact equivalent to $\min_{\mathbf{x} \in K} \theta_{\text{gap}}(\mathbf{x})$. To see why, let $K_D(\mathbf{x})$ denote the feasible region of $DLP(\mathbf{x})$, as given by (3.7). Indeed, since K is closed,

$$\begin{aligned} \theta_{\text{gap}}(\mathbf{x}) &= \mathbf{F}(\mathbf{x})'\mathbf{x} - \min_{\mathbf{y} \in K} \mathbf{F}(\mathbf{x})'\mathbf{y} \\ &= \mathbf{F}(\mathbf{x})'\mathbf{x} - \max_{\boldsymbol{\lambda} \in K_D(\mathbf{x})} \mathbf{b}'\boldsymbol{\lambda} \\ &= \min_{\boldsymbol{\lambda} \in K_D(\mathbf{x})} \left[\mathbf{F}(\mathbf{x})'\mathbf{x} - \mathbf{b}'\boldsymbol{\lambda}\right] . \end{aligned}$$

Recall that the objective and constraint functions of reformulation (3.8) are continuously differentiable as many times as $\mathbf{F}(\mathbf{x})$ is. Accordingly, the classical primal gap function does in fact induce a continuously differentiable merit-type function, but it is one that is defined over a modestly higher-dimensional space.

3.3 Concise Reformulations

Inspired by the robust optimization paradigm, our results in the previous section extend the equivalence with single-level optimization to a larger class of VIs. One may note that this extension appears to come at a minor cost, in that our reformulations are in an augmented space of higher dimension than the original problem. We underscore that this enlargement is indeed modest in size; whereas K, as given in (3.4), is defined by $m + n$ constraints in n variables, the VI reformulation (3.8) is defined by $m + 2n$ constraints in $m + n$ variables.

Moreover, as we now show, we can in fact give a reduced, i.e., a more concise, reformulation of $VI(K, \mathbf{F})$, by eliminating redundant variables and constraints from (3.8). This more concise reformulation preserves the convexity of (3.8).

Just as we restricted our attention, without loss of generality to the case of K in standard form, we may furthermore assume, without loss of generality, that $\mathbf{A}$, as it appears in definition (3.4) of K, is of full rank (Theorem 2.5 of Bertsimas and Tsitsiklis [19]). Accordingly, in the optimization reformulation (3.8) of $VI(K, \mathbf{F})$, we may use any basis matrix $\mathbf{B}$ of $\mathbf{A}$, together with the constraint $\mathbf{Ax} = \mathbf{b}$, to eliminate redundant variables. In particular, without loss of generality, assume that $\mathbf{A} = [\mathbf{B}, \mathbf{A}_N]$, where $\mathbf{A}_N$ are the non-basic columns of $\mathbf{A}$ induced by the basis matrix $\mathbf{B}$. Let the corresponding vectors of basic and nonbasic variables be denoted by $\mathbf{x}_B$ and $\mathbf{x}_N$, respectively. Let

$$\begin{aligned} \mathbf{Q} &= \begin{pmatrix} -\mathbf{B}^{-1}\mathbf{A}_N \\ \mathbf{I} \end{pmatrix} \in \mathbb{R}^{n \times (n-m)} \\ \mathbf{r} &= \begin{pmatrix} \mathbf{B}^{-1}\mathbf{b} \\ \mathbf{0} \end{pmatrix} \in \mathbb{R}^{n}, \end{aligned}$$

where $\mathbf{I}$ and $\mathbf{0}$ are the appropriately dimensioned identity matrix and vector of zeros, respectively. We can thus eliminate $\mathbf{x}_B$ by rewriting

$$\begin{pmatrix} \mathbf{x}_B \\ \mathbf{x}_N \end{pmatrix} = \mathbf{Q}\mathbf{x}_N + \mathbf{r}.$$

Consequently, the optimization reformulation (3.8) of $VI(K, \mathbf{F})$ can equivalently be rewritten as the following MP in $2n$ constraints and n variables.

$$\begin{aligned} \min_{\mathbf{x}_N, \boldsymbol{\lambda}} \quad & \mathbf{F}(\mathbf{Q}\mathbf{x}_N + \mathbf{r})' [\mathbf{Q}\mathbf{x}_N + \mathbf{r}] - \mathbf{b}'\boldsymbol{\lambda} \\ & \mathbf{Q}\mathbf{x}_N + \mathbf{r} \geq \mathbf{0} \\ & \mathbf{A}'\boldsymbol{\lambda} \leq \mathbf{F}(\mathbf{Q}\mathbf{x}_N + \mathbf{r}). \end{aligned} \tag{3.15}$$

Furthermore, it is clear that convexity of formulation (3.8) implies that of (3.15), since, in moving from the former to the latter, we have simply projected the former onto a lower-dimensional space.

An alternative concise formulation may be given in the case of affine $VI(K, \mathbf{F})$, with $\mathbf{F}(\mathbf{x}) = \mathbf{G}\mathbf{x} + \mathbf{h}$, and $\mathbf{G}$ nonsingular. For this purpose, let us assume that K is not in standard form, but instead given by

$$K = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{A}\mathbf{x} \geq \mathbf{b}\} \neq \emptyset,$$

where $\mathbf{A} \in \mathbb{R}^{\ell \times n}$ and $\mathbf{b} \in \mathbb{R}^{\ell}$. This form of K yields a slightly different form of optimization problem than (3.8) that is equivalent to $VI(K, \mathbf{F})$. Namely, we obtain the following MP in $n + 2\ell$ constraints and $n + \ell$ variables.

$$\begin{aligned} \min_{\mathbf{x}, \boldsymbol{\lambda}} \quad & \mathbf{x}'\mathbf{G}\mathbf{x} + \mathbf{h}'\mathbf{x} - \mathbf{b}'\boldsymbol{\lambda} \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} \geq \mathbf{b} \\ & \mathbf{A}'\boldsymbol{\lambda} = \mathbf{G}\mathbf{x} + \mathbf{h} \\ & \boldsymbol{\lambda} \geq \mathbf{0}. \end{aligned} \tag{3.16}$$

We may eliminate $\mathbf{x}$ by rewriting it in terms of $\boldsymbol{\lambda}$.

$$\mathbf{x} = \mathbf{G}^{-1}\mathbf{A}'\boldsymbol{\lambda} - \mathbf{G}^{-1}\mathbf{h}.$$

Letting

$$\begin{aligned}
\hat{\mathbf{G}} &= \mathbf{A}\left[\mathbf{G}^{-1}\right]'\mathbf{A}' \\
\hat{\mathbf{h}} &= -\mathbf{h}'\left[\mathbf{G}^{-1}\right]'\mathbf{A}' \\
\hat{\mathbf{A}} &= \mathbf{A}\mathbf{G}^{-1}\mathbf{A}' \\
\hat{\mathbf{b}} &= \mathbf{A}\mathbf{G}^{-1}\mathbf{h} + \mathbf{b},
\end{aligned}$$

formulation (3.16) then becomes the following MP in 2ℓ constraints and ℓ variables.

$$\begin{aligned}
\min_{\boldsymbol{\lambda}} \quad & \boldsymbol{\lambda}'\hat{\mathbf{G}}\boldsymbol{\lambda} + \left(\hat{\mathbf{h}} - \mathbf{b}\right)'\boldsymbol{\lambda} \\
\text{s.t.} \quad & \hat{\mathbf{A}}\boldsymbol{\lambda} \geq \hat{\mathbf{b}} \qquad (3.17) \\
& \boldsymbol{\lambda} \geq \mathbf{0}.
\end{aligned}$$

Again, it is clear that convexity of formulation (3.16) implies that of (3.17), since $\mathbf{G} \succeq \mathbf{0}$ implies $\hat{\mathbf{G}} \succeq \mathbf{0}$.

Chapter 4

Robust Game Theory

4.1 Introduction

4.1.1 Finite Games with Complete Information

Game theory is a field in economics that examines multi-agent decision problems, in which the rewards to each agent, or player, can depend not only on his action, but also on the actions of the other players. In his seminal paper [116], John Nash introduced the notion of an equilibrium of a game. He defined an equilibrium as a profile of players' strategies, such that no player has incentive to unilaterally deviate from his strategy, given the strategies of the other players.

In [116] and [117], Nash focused on non-cooperative, simultaneous-move, one-shot, finite games with complete information, a class of games encompassing various situations in economics. "Simultaneous-move" refers to the fact that the players choose strategies without knowing those selected by the other players. "One-shot" means that the game is played only once. "Finite" connotes that there are a finite number of players, each having a finite number of actions, over which mixed strategies in these actions may be defined. Finally, "complete information" implies that all parameters of the game, including individual players' payoff functions, are common knowledge.

In his analysis, Nash modeled each player as rational and wanting to maximize his expected payoff with respect to the probability distributions given by the mixed strategies. Nash proved that each game of the aforementioned type has an equilibrium in mixed strategies. In fact, Nash gave two existence proofs, one in [116] based on Kakutani's Fixed Point Theorem [79] and one in [117] based on the less general Brouwer's Fixed Point Theorem [28].

Nash's equilibrium concept and existence theorem have become a cornerstone in the field of game theory and earned him the 1994 Nobel Prize in Economics. The concept is regarded as practically significant largely because, under the standard assumptions that players are rational and that the structure of a game is common knowledge, the concept offers a possible approach to predicting the outcome of the game. The argument is as follows. Any rational player who thinks his opponents will use certain strategies should never play anything other than a best response to those strategies. By the common knowledge assumption, the other players know this, the player knows that the other players know this, etc., *ad infinitum.* Thus, the players may be able to reach consistent predictions of what each other will play. The classical game theory literature concludes from this observation that we should expect the realized behavior in a game to belong to the set of equilibria. As discussed in the introduction of Fudenberg and Levine [53], in practice, this conclusion can prove to be unreliable. Nonetheless, the concept of Nash equilibrium remains the central idea in game theory, in part because no solution concept has been offered that overcomes these prediction issues.

4.1.2 Finite Games with Incomplete Information

While the existence of an equilibrium can be asked in any game, Nash's existence result addresses only non-cooperative, simultaneous-move, one-shot, finite games with complete information. Of course, in real-world, game-theoretic situations, players are often uncertain of some aspects of the structure of the game, such as payoff functions.

Harsanyi [66] modeled these incomplete-information games as what he called "Bayesian games." He defined a player's "type" as an encoding of information available to that player, including knowledge of his own payoff function and beliefs about other players' payoff functions. In this way, he used type spaces to model incomplete-information games, in which some players may have private information. He assumed that the players share a common-knowledge prior probability distribution over the type space. Harsanyi suggested that each player would use this prior probability distribution, together with his type, to derive a conditional probability distribution on the parameter values remaining unknown to him. Furthermore, he assumed that each player's goal would then be to, using a Bayesian approach in this way, maximize his expected payoff with respect to both the aforementioned conditional probability distribution and the mixed strategies of the players.

In this framework, Harsanyi extended Nash's result to non-cooperative, simultaneous-move, one-shot, finite games with *incomplete* information. In particular, he showed that any Bayesian game is equivalent to an extensive-form game with complete, but imperfect information. This extensive-form game, in turn, is known to have a static-form representation. Using an equilibrium existence result more general than Nash's and due to Debreu [41], Harsanyi thus proved the existence of equilibria, which he called "Bayesian equilibria," in Bayesian games. For his work on games with incomplete information, he won the 1994 Nobel Prize in Economics, alongside Nash and Selten.

Harsanyi's work relaxes the assumption that all parameters affecting the payoffs of the players are known with certainty. His model technique is essentially analogous to the stochastic programming approach to data uncertainty in mathematical optimization problems. As in stochastic programming, Harsanyi's model assumes the availability of full prior distributional information for all unknown parameters. In addition, his analysis assumes that all players use the same prior, and that this fact is common knowledge. Many, including Morris [112] and Wilson [165], have questioned the common prior and common knowledge aspects of these assumptions. Nonetheless, Harsanyi's Bayesian model remains the accepted convention for analyzing static games with incomplete information.

Some contributions to the literature have relaxed the common prior and common knowledge assumptions of Harsanyi's model. Perhaps most importantly, Mertens and Zamir [109] formalized the notion of a "universal type space," a type space large enough to capture players' higher-order beliefs, players' use of different prior probability distributions on the uncertain parameters, and the absence of common knowledge of these priors.

Taking a different approach, other contributors to the game theory literature have offered distribution-free equilibrium concepts for incomplete-information games. These analyses address the possibility that distributional information is not available to the players, or that they opt not to use potentially inaccurate distributional information. The notion of an *ex post* equilibrium is the most common distribution-free solution concept and is especially prevalent in the auction theory literature. Holmström and Myerson [75] first introduced this notion under the name "uniform incentive compatibility," and Crémer and McLean [37] first used it in the context of auctions. The *ex post* equilibrium is a refinement of the Bayesian equilibrium and is an appealing solution concept, because it is relevant even when the players lack distributional information on the uncertain parameters. However, it is a strong concept, and *ex post* equilibria need not exist in an incomplete-information game.

4.1.3 Contributions and Structure of this Chapter

The contributions of this chapter of the thesis are as follows.

1. We propose a distribution-free, robust optimization model for games with incomplete information. Our model relaxes the assumptions of Harsanyi's Bayesian games model, and at the same time provides a more general equilibrium concept than that of the *ex post* equilibrium.

   Specifically, in Section 4.2, we formally propose the robust optimization model for non-cooperative, simultaneous-move, one-shot, finite games with incomplete information. We start by discussing precedents, from the game theory literature, for using a worst-case approach to uncertainty, in the absence of probability distributions. We then note the novelty of our approach with respect to these works. After setting forth our model, we describe the "robust games," analogous to Harsanyi's Bayesian games, to which this approach gives rise. We compare the equilibrium conditions for such robust games to those for Bayesian equilibria and *ex post* equilibria. Furthermore, we note that any *ex post* equilibria of an incomplete-information game are what we call "robust-optimization equilibria," i.e., equilibria of the corresponding robust game, just as they are Bayesian equilibria under Harsanyi's model. We then discuss our union of the notion of equilibrium with the robust optimization paradigm, and we give interpretations of mixed strategies in the context of robust games. We relate this discussion and these interpretations to those in the literature on complete-information games.

At the end of Section 4.2, to concretize the idea of a robust game, we present some examples. In addition, we use one of these examples to illustrate that *ex post* equilibria need not exist, thereby motivating the need for an alternate distribution-free equilibrium concept.

Let us pause to note that, for the sake of simplicity, in Sections 4.2 through 4.5, we focus on situations of uncertainty in which no player has private information. This focus allows for a clearer and a sufficiently rich discussion of the main ideas underlying our model and results, without hindering the reader's understanding through the use of cumbersome notation and references to results from the theory of Banach spaces. Such notation and results are required for the general case of incomplete-information games involving potentially private information. In Section 4.6, we extend our analysis to this general case.

2. In Section 4.3, we prove the existence of equilibria in robust finite games with bounded uncertainty sets and no private information.

3. In Section 4.4, we formulate the set of equilibria of an arbitrary robust finite game, with bounded polyhedral uncertainty set and no private information, as the dimension-reducing, component-wise projection of the solution set of a system of multilinear equalities and inequalities. We provide a general method for approximately computing sample robust-optimization equilibria, and we present numerical results from the application of this method. For a special class of such games, we show equivalence to finite games with complete payoff information and the same action spaces. As a result, in order to compute sample equilibria for robust games in this class, one need only

solve for the equilibria of the corresponding complete-information game.

4. In Section 4.5, we compare properties of robust finite games with those of the corresponding complete-information games, in which the uncertain payoff parameters of the former are commonly known to take fixed, nominal values. In the absence of private information, these nominal games are precisely the Bayesian games arising from attributing symmetric probability distributions over the uncertainty sets of the corresponding robust games. In addition, turning our attention to a notion of symmetry unrelated to the symmetry of probability distributions, we extend the definition of a symmetric game, i.e., one in which the players are indistinguishable with respect to the structure of the game, to the robust game setting. We prove the existence of symmetric equilibria in symmetric, robust finite games with bounded uncertainty sets and no private information.

5. In Section 4.6, we generalize our model to the case with potentially private information. We extend our existence result to this context and generalize our computation method to such situations involving private information and finite type spaces.

4.1.4 Notation

In addition to the notation conventions outlined in Section 1.3, we will also use the following shorthand. $\text{vec}(\mathbf{A})$ will denote the column vector obtained by stacking the row vectors of the matrix $\mathbf{A}$ one on top of another. Thus, if $\mathbf{A}$ is an $m \times n$ matrix, $\text{vec}(\mathbf{A})$ will be a $mn \times 1$ vector.

4.2 A Robust Approach to Payoff Uncertainty in Games

As an alternative to Harsanyi's model and the notion of the *ex post* equilibrium, we propose a new distribution-free model of and equilibrium concept for incomplete-information games. Our model is based on robust optimization, in which one takes a deterministic approach to uncertainty and seeks to optimize worst-case performance, where the worst case is taken with respect to a set of possible values for the uncertain parameters. Let us note that, in Sections 4.2 through 4.5, we will focus on incomplete-information games without private information. In Section 4.6, we will extend our analysis to the general case involving potentially private information.

4.2.1 Precedents for a Worst-Case Approach

In fact, the game theory literature is ripe with precedents for using a worst-case approach. The field arose in large part from von Neumann's and Morgenstern's "max-min" formulation of behavior in games [161]. More recently, for example, Goldberg et al. [59] proposed a worst-case, competitive-analysis approach to auction mechanism design. In the more general context of normal-form games, several authors, including Gilboa and Schmeidler [58], Dow and Werlang [43], Klibanoff [81], Lo [94], and Marinacci [106], have argued for a max-min-based approach to "ambiguous uncertainty," uncertainty in the absence of probabilistic information. They contend that expected utility models are well-suited for decision-theoretic situations characterized by "risk," uncertainty with distributional information, but that these models do not capture behavior observed in practice in situations of ambiguous uncertainty.[1] As Dow and Werlang [43] note, the former type of uncertainty is exemplified by the outcome of a coin toss, while the latter is typified by the outcome of a horse race.

While the max-min approaches of Gilboa and Schmeidler, Dow and Werlang, Klibanoff, Lo, and Marinacci share the worst-case perspective of our model, their approaches are fundamentally different from ours for at least three reasons. First, these authors consider complete-information games, whereas we address incomplete-information games. In their models, players know, with certainty, the payoffs under given tuples of actions, but do not know which tuple of actions will be played. In our model, the players may be uncertain of the payoffs, even under given tuples of actions. Accordingly, the aforementioned authors use a pessimistic approach to model each player's uncertainty of the other players' behaviors, whereas we use a worst-case approach to model each player's uncertainty of the payoff

[1]Knight [82] was one of the first to draw a distinction between these two forms of uncertainty.

functions themselves.

Second, although these authors take a worst-case approach to some extent, their models are nonetheless inherently probabilistic, unlike our approach, which is fundamentally deterministic. Klibanoff [81] and Lo [94] model each player's uncertainty of the other players' behaviors using the notion of multiple prior probability distributions. They characterize each player as believing his counterparts' actions are a realization from some unknown probability distribution, belonging to a family of known multiple priors. Each player then seeks to maximize his minimum expected utility, where the minimum is taken with respect to this set of multiple priors. Gilboa and Schmeidler [58], Dow and Werlang [43], and Marinacci [106] propose a related approach using non-additive probability distributions in place of sets of multiple priors. Unlike these authors, we offer a model in which the players give no consideration whatsoever to probability distributions over the uncertain values. Under our approach, the players regard the uncertain values as simply unknown and not as realizations from some probability distribution, even a distribution that is itself not exactly known. Consequently, our model of the players' responses to uncertainty is distribution-free and deterministic in nature.

Third, these authors offer no guidelines for equilibria computation in the context of their models. In contrast, in Section 4.4, we propose such a computation method. Despite these differences, the aforementioned authors' contributions provide ample support for the robust optimization model we propose for games with incomplete information.

In addition, since the submission of this chapter of the thesis for publication, Hyafil and Boutilier [77] have recently offered a worst-case approach for incomplete-information games, based on the distribution-free decision criterion of minimax regret, popular in the online optimization literature [24]. Their approach is in contrast to our framework of modeling the players as each seeking to maximize his worst-case expected payoff. Hyafil and Boutilier provide an existence result for a very restricted, special case of incomplete-information games, involving private information, but finite type spaces. They offer no ideas on computation of their equilibria.

Having discussed some precedents for taking a worst-case approach to analyzing game-theoretic situations, let us now formalize our robust games model.

4.2.2 Formalization of the Robust Game Model

In our robust optimization model of incomplete-information games, we assume that the players commonly know only an uncertainty set of possible values of the uncertain payoff function parameters.[2] They need not, as Harsanyi's model additionally assumes, have distributional information for this uncertainty set. In addition, we suppose that each player uses a robust optimization, and therefore a worst-case, approach to the uncertainty, rather than seeking, as in Harsanyi's model, to optimize "average" performance with respect to a distribution over the uncertainty set. In the game theory literature, the "performance" of a player's mixed strategy is measured by his expected payoff. Accordingly, in our model, given the other players' strategies, each player seeks to maximize his worst-case expected payoff. The worst-case is taken with respect to the uncertainty set, and the expectation is taken, as in complete-information games, over the mixed strategies of the players. Analogous to Harsanyi's "Bayesian game" terminology, we call the resulting games "robust games," and we refer to their equilibria as "robust-optimization equilibria" of the corresponding incomplete-information games.

In this section we will formalize our robust game model and its relation to Nash's and Harsanyi's models for the complete- and incomplete-information settings, respectively. We will also compare the notion of *ex post* equilibrium with the concept of robust-optimization equilibrium.

Let us first define some terms and establish some notation. Suppose there are N players and that player $i \in \{1, \dots, N\}$ has $a_i > 1$ possible actions.

Definition 4.2.1. A game is said to be **finite** if the number of players N and the number of actions a_i available to each player $i \in \{1, \dots, N\}$ are all finite.

[2] Incomplete-information games in the absence of distributional information are sometimes called "games in informational form" [76].

So, we will use the term "robust finite game" to refer to robust games that have finitely many players with finitely many actions each, even when the uncertainty sets are not finite.

In the complete-information game setting, a multi-dimensional payoff matrix $\check{\mathbf{P}}$, indexed over $\{1,\ldots,N\} \times \prod_{i=1}^{N}\{1,\ldots,a_i\}$, records the payoffs to the players under all possible action profiles for the players. In particular, for $i \in \{1,\ldots,N\}$, $(j_1,\ldots,j_N) \in \prod_{i=1}^{N}\{1,\ldots,a_i\}$, let $\check{P}^i_{(j_1,\ldots,j_N)}$ denote the payoff to player i when player $i' \in \{1,\ldots,N\}$ plays action $j_{i'} \in \{1,\ldots,a_{i'}\}$. Let

$$S_{a_i} \;=\; \left\{ \mathbf{x}^i \in \mathbb{R}^{a_i} \;\middle|\; \mathbf{x}^i \geq 0,\; \sum_{j_i=1}^{a_i} x^i_{j_i} = 1 \right\}.$$

That is, S_{a_i} is the set of mixed strategies over action space $\{1,\ldots,a_i\}$. Let us define $\boldsymbol{\pi} : U \times \prod_{i=1}^{N} S_{a_i} \to \mathbb{R}^N$ as the vector function mapping a payoff matrix and the mixed strategies of N players to a vector of expected payoffs to the N players. In particular, $\pi_i\left(\mathbf{P};\mathbf{x}^1,\ldots,\mathbf{x}^N\right)$ will denote the expected payoff to player i when the payoff matrix is given by $\mathbf{P}$ and player $i' \in \{1,\ldots,N\}$ plays mixed strategy $\mathbf{x}^{i'} \in S_{a_{i'}}$. That is,

$$\pi_i\left(\mathbf{P};\mathbf{x}^1,\ldots,\mathbf{x}^N\right) \;=\; \sum_{j_1=1}^{a_1}\cdots\sum_{j_i=1}^{a_i}\cdots\sum_{j_N=1}^{a_N} P^i_{(j_1,\ldots,j_N)} \prod_{i=1}^{N} x^i_{j_i}.$$

Now that we have established some notation, we can formulate the best response correspondence in our robust optimization model for games with incomplete payoff information. We will compare this correspondence with those in Nash's and Harsanyi's models for the complete- and incomplete-information cases, respectively. In the remainder of this chapter, we will use the following shorthands.

$$\mathbf{x}^{-i} \triangleq \left(\mathbf{x}^1,\ldots,\mathbf{x}^{i-1},\mathbf{x}^{i+1},\ldots,\mathbf{x}^N\right), \qquad S \triangleq \prod_{i=1}^{N} S_{a_i}, \tag{4.1}$$

$$\left(\mathbf{x}^{-i},\mathbf{u}^i\right) \triangleq \left(\mathbf{x}^1,\ldots,\mathbf{x}^{i-1},\mathbf{u}^i,\mathbf{x}^{i+1},\ldots,\mathbf{x}^N\right), \qquad S_{-i} \triangleq \prod_{\substack{i'=1\\ i'\neq i}}^{N} S_{a_{i'}}.$$

Every model of a game attributes some objective to each player. A player's objective in turn determines the set of best responses to the other players' strategies.

Definition 4.2.2. A player's strategy is called a **best response** to the other players' strategies if, given the latter, he has no incentive to unilaterally deviate from his aforementioned strategy.

In the complete-information game setting, with payoff matrix $\check{\mathbf{P}}$, the classical model assumes that each player seeks to maximize his expected payoff. So, player i's best response to the other players' strategies $\mathbf{x}^{-i} \in S_{-i}$ belongs, by definition, to

$$\arg\max_{\mathbf{u}^i \in S_{a_i}} \pi_i \left(\check{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right).$$

In games with incomplete payoff information, the payoff matrix $\tilde{\mathbf{P}}$ is subject to uncertainty. In Harsanyi's Bayesian model, in the context of games without private information, in which the type spaces are singletons, player i's best response to the other players' strategies $\mathbf{x}^{-i} \in S_{-i}$ must belong to

$$\arg\max_{\mathbf{u}^i \in S_{a_i}} \left[\mathop{\mathrm{E}}_{\tilde{\mathbf{P}}} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right)\right].$$

In our robust model, for the case without private information, player i's best response to the other players' strategies $\mathbf{x}^{-i} \in S_{-i}$ must belong to

$$\arg\max_{\mathbf{u}^i \in S_{a_i}} \left[\inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right)\right].$$

Thus, in moving from Harsanyi's Bayesian approach to our robust optimization model, we replace the expectation in the definition of the best response correspondence with an infimum operator.

Note that $\forall i \in \{1, \ldots, N\}$ and $\forall \left(\mathbf{x}^{-i}, \mathbf{u}^i\right) \in S$, by the linearity of π_i over U and by the linearity of the expectation operator,

$$\mathop{\mathrm{E}}_{\tilde{\mathbf{P}}} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right) = \pi_i \left(\mathop{\mathrm{E}}_{\tilde{\mathbf{P}}} \left[\tilde{\mathbf{P}}\right]; \mathbf{x}^{-i}, \mathbf{u}^i\right), \tag{4.2}$$

where $\mathop{\mathrm{E}}_{\tilde{\mathbf{P}}} \left[\tilde{\mathbf{P}}\right]$ denotes the component-wise expected value of $\tilde{\mathbf{P}}$. Hence, in the Bayesian game setting, the average expected payoffs and expected average payoffs are in fact equivalent.[3] Recall, from Harsanyi [66], that any Bayesian game with incomplete information is equivalent to a static game with complete but imperfect information. As indicated by Equation (4.2), in the absence of private information, a Bayesian game is equivalent to a finite game with complete and *perfect* information, with the same action spaces and with payoff matrix $\mathop{\mathrm{E}}_{\tilde{\mathbf{P}}} \left[\tilde{\mathbf{P}}\right]$.

In contrast, under the robust model, the worst-case expected payoff expressed above is no less than, and is generally strictly greater than, the expected worst-case payoff. That is,

$$\inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right) \geq \pi_i \left(\inf_{\tilde{\mathbf{P}} \in U} \left[\tilde{\mathbf{P}}\right]; \mathbf{x}^{-i}, \mathbf{u}^i\right),$$

where $\inf_{\tilde{\mathbf{P}} \in U} \left[\tilde{\mathbf{P}}\right]$ denotes the component-wise infimum of $\tilde{\mathbf{P}}$. Thus, a robust finite game without private information is, in general, not equivalent to the complete-information, finite game with the same action spaces and with payoff matrix commonly known to be $\inf_{\tilde{\mathbf{P}} \in U} \left[\tilde{\mathbf{P}}\right]$. We will see in Section 4.4 that this equivalence does, however, hold for certain classes of robust games.[4]

[3] We use the terms "average" and "expected" in an effort to distinguish between two different types of expectations, namely, the expectation ("average") taken with respect to the distribution over the uncertainty set of payoff parameter values and the expectation ("expected payoff") taken with respect to the distributions induced by the players' mixed strategies.

[4] One could model each player as wishing to maximize his expected worst-case payoff, rather than, as we have done, his worst-case expected payoff. We chose the latter over the former for two reasons. First, the former model is *not*, while the latter model is, in the spirit of robust optimization, in which we seek to optimize a worst-case version of the nominal objective, i.e., the expected payoff. Second, while a robust approach is by its nature pessimistic, the former model is even more, and perhaps excessively, pessimistic. In it, each player assumes that the uncertain data realization will be maximally hostile with respect to the *action outcomes* of the randomizations yielded by all the players' mixed strategies. In contrast, in the robust model we propose, the maximal hostility is assumed by the players to be with respect to the mixed strategy *probability distributions* themselves; i.e., the "adversary" does not have the benefit of seeing the action outcomes of the randomizations, before he is forced to choose values of the uncertain data.

If one nonetheless opts, despite these drawbacks, to model each player as seeking to maximize his expected

We are now ready to apply the concept of equilibrium to robust finite games.

Definition 4.2.3. A tuple of strategies is said to be an **equilibrium** if each player's strategy is a best response to the other players' strategies.

Accordingly, the criterion for an equilibrium is completely determined by the best response correspondence, which in turn is completely determined by the players' objectives. For example, in the complete-information game setting, $\left(\mathbf{x}^1, \dots, \mathbf{x}^N\right) \in S$ is said to be a Nash equilibrium iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{x}^i \in \arg\max_{\mathbf{u}^i \in S_{a_i}} \pi_i\left(\check{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right). \tag{4.3}$$

Similarly, under Harsanyi's model for finite games with incomplete payoff information and with no private information, $\left(\mathbf{x}^1, \dots, \mathbf{x}^N\right) \in S$ is said to be an equilibrium iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{x}^i \in \arg\max_{\mathbf{u}^i \in S_{a_i}} \left[\operatorname*{E}_{\tilde{\mathbf{P}}} \pi_i\left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right)\right]. \tag{4.4}$$

Finally, under our robust model for finite games with incomplete payoff information and with no private information, $\left(\mathbf{x}^1, \dots, \mathbf{x}^N\right) \in S$ is said to be an equilibrium, i.e., a robust-optimization equilibrium of the corresponding game with incomplete information, iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{x}^i \in \arg\max_{\mathbf{u}^i \in S_{a_i}} \left[\inf_{\tilde{\mathbf{P}} \in U} \pi_i\left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right)\right]. \tag{4.5}$$

Let us contrast the equilibrium concepts arising from Harsanyi's Bayesian game model and our robust game model with the notion of the *ex post* equilibrium, defined as follows.

Definition 4.2.4. A tuple of strategies is said to be an **ex post equilibrium** if each player's strategy is a best response to the other players' strategies, under all possible realizations of

worst-case payoff, rather than his worst-case expected payoff, the game with incomplete information will be equivalent to one with complete information and with payoff matrix $\inf_{\tilde{\mathbf{P}} \in U}\left[\tilde{\mathbf{P}}\right]$. Accordingly, the existence and computation results that we will present in this chapter for the robust model follow trivially for this excessively pessimistic model.

the uncertain data.

More precisely, in the absence of private information, $\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right) \in S$ is said to be an *ex post* equilibrium iff, $\forall i \in \{1, \ldots, N\}$,

$$\mathbf{x}^i \in \arg\max_{\mathbf{u}^i \in S_{a_i}} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right), \qquad \forall \tilde{\mathbf{P}} \in U. \tag{4.6}$$

By definition, an *ex post* equilibrium must be an equilibrium of every nominal game in the family of nominal games arising from U. This condition is quite strong. In fact, it is easy to show that every *ex post* equilibrium of an incomplete-information game is an equilibrium of any corresponding Bayesian game arising from the assignment of a distribution over the set U. Similarly, we have the following lemma, establishing an analogous result for the set of robust-optimization equilibria.

Lemma 4.2.1. *The set of ex post equilibria of an incomplete-information game with no private information is contained in the corresponding set of robust-optimization equilibria.*

Proof. Suppose $\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right) \in S$ is an *ex post* equilibrium of the incomplete-information game with uncertainty set U. Suppose, $\exists i \in \{1, \ldots, N\}$ and $\exists \mathbf{u}^i \in S_{a_i}$, such that

$$\inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i\right) < \inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right).$$

By the definition of *ex post* equilibrium,

$$\pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right) \leq \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i\right), \qquad \forall \tilde{\mathbf{P}} \in U,$$

yielding a contradiction of the fact that $\inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i\right)$ is the greatest lower bound on $\pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i\right)$ over $\tilde{\mathbf{P}} \in U$. Therefore, $\forall i \in \{1, \ldots, N\}$, and $\forall \mathbf{u}^i \in S_{a_i}$,

$$\inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i\right) \geq \inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right),$$

establishing that $\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right) \in S$ is an equilibrium of the corresponding robust game. □

In Section 4.2.5, we will illustrate, with examples, our robust optimization model for games with incomplete information. Using one of these examples, in Section 4.2.6, we will demonstrate that, in general, *ex post* equilibria do not exist in incomplete-information games. Before giving these examples, we wish to address two questions that the reader may have regarding our approach. In Section 4.2.3, we will discuss why, in the context of distribution-free, incomplete-information games, it is reasonable, and in fact natural, to combine the notion of equilibrium with a worst-case viewpoint. In Section 4.2.4, we will discuss our motivation for considering mixed strategies.

4.2.3 Why Combine Equilibrium and Worst-Case Notions?

Recall that, with the exception of two-person, zero-sum games with complete information, mixed strategy equilibria do not generally consist of max-min strategies. That is, a player's strategy in a mixed-strategy equilibrium is not generally the one guaranteeing him the best possible expected payoff when his counterparts collude to minimize this quantity. The reason is that a player's counterparts generally have incentive to deviate from such collusive behavior, in order to try to individually maximize their own payoffs. In turn, the player himself therefore generally has incentive to deviate from the aforementioned max-min strategy.

In contrast, the robust optimization paradigm is fundamentally such a max-min, or a worst-case, approach. In our robust games model, given his counterparts' strategies, each player formulates a best response as the solution of a robust optimization problem. Based on the discussion in the preceding paragraph, one may worry that, by analogy, best responses based on robust optimization are not conducive to equilibrium. This analogy fails, and this worry is therefore unfounded, for the following reason. In our model, a player's counterparts are outside the scope of that player's pessimistic viewpoint. In particular, each player takes a worst-case view *only* of the uncertain parameters that define his payoff function, under a *given* tuple of his counterparts' strategies. Each player does *not* take a worst-case approach to his uncertainty with respect to this tuple itself, as is done in classical max-min strategies. Indeed, "nature," rather than any of the players themselves, selects these unknown payoff parameter values. Accordingly, in order for the analogy to hold, nature must be a participant, on the same footing as the other players, in the game. However, nature receives no payoff in the game, and therefore cannot be characterized as a player itself.

Thus, it is indeed reasonable to combine, as we have done, the notion of equilibrium with the robust optimization paradigm. Let us now explain why this union is in fact natural, in the context of incomplete-information games. If the players commonly know that they all take a robust optimization approach to the payoff uncertainty, then they would all commonly know each others' best response correspondences. Armed with this common knowledge, the players could then attempt to mutually predict each other's behavior, just as they could in a complete-information game, as discussed in Section 4.1.1. Recall from this discussion that the set of Nash equilibria are the set of *consistent* such mutual predictions in a finite, complete-information game. Analogously, the set of equilibria of a robust finite game are the set of consistent such mutual predictions in the corresponding finite, distribution-free, incomplete-information game. As such, our notion of equilibrium in a robust game offers a natural approach to attempting to predict the outcomes of such incomplete-information games.

4.2.4 Interpretation of Mixed Strategies

We will now explain our motivation for considering mixed strategies, and we will relate this discussion to interpretations of mixed strategies in the context of complete-information games (see, for example, Chapter 3 of Osborne and Rubinstein [119]). In the case of finite games with complete information, some game theorists support the literal interpretation of mixed strategies as actual randomizations by the players over their action spaces. Others are dissatisfied with this viewpoint. The latter group note the following property of mixed strategy equilibria, in finite, complete-information games. In response to his counterparts' behaviors in any such equilibrium, each player's mixed strategy does as well, but no better than the actions contained in its support. The opponents of the literal interpretation therefore argue that this lack of strict preference for randomization undermines the belief that players randomize in reality.

In the case of robust finite games, this argument against the literal interpretation does not hold. In particular, because of the infimum in the worst-case expected payoff function, this function is nonlinear. Consequently, for any mixed strategy equilibrium in such a game, in response to his counterparts' behavior in this equilibrium, each player will, in general, strictly prefer his mixed strategy over the actions in its support. Accordingly, one may argue that the literal viewpoint of mixed strategies is more justified in the context of robust games than it is in the context of complete-information games.

One may nonetheless remain dissatisfied with this belief that players randomize in real-world, game-theoretic settings, even those involving payoff uncertainty. Let us then consider an alternative interpretation of mixed strategies. In the literature on finite, complete-information games, some have advocated the viewpoint of mixed strategy equilibria as limiting, empirical frequencies of actions played, when the game is repeated.

The same empirical frequency interpretation extends to robust finite games. Imagine that the players engage concurrently in many instances of the same game, with the same, unknown payoff matrix $\tilde{\mathbf{P}}$. Suppose the players know that $\tilde{\mathbf{P}}$ is constant across all instances, but are uncertain of its true value. As before, suppose each player knows only an uncertainty set to which $\tilde{\mathbf{P}}$ belongs, has no distributional information with respect to this set, and takes a worst-case approach to this uncertainty. Lastly, suppose that, in each instance of the game, each agent may play a different action. Each player thus builds, in essence, a "portfolio" of actions. The payoff from each action in the portfolio is determined by the other players' actions in the corresponding instance of the game and by the single unknown value of $\tilde{\mathbf{P}}$. Accordingly, we may view the mixed strategy equilibria as the limiting, empirical frequencies describing each player's level of diversification within his portfolio of actions. Note that this portfolio interpretation can be recast in terms of sequentially repeated games, in which the players know that the uncertain payoff matrix is constant over all rounds, and in which they do not receive their payoffs until the final round is played. That is, the players do not know, until at least after play has terminated, the true value of $\tilde{\mathbf{P}}$.

4.2.5 Examples of Robust Finite Games

Having presented our robust games model and addressed some interpretation issues, we will now illustrate our approach with a few examples.

Example 1. Robust Inspection Game

Consider the classical inspection game discussed in [54]. The row player, the employee, can either shirk or work (actions 1 and 2, respectively). The column player, his employer, can either inspect or not inspect (actions 1 and 2, respectively). The purpose of inspecting is to learn whether the employee is working. The two players simultaneously select their actions. When the employee works, he suffers an opportunity cost $\tilde{g}$, and his employer enjoys a value of work output of $\tilde{v}$. When the employer inspects, she suffers an opportunity cost of $\tilde{h}$. If she inspects and finds the employee shirking, she need not pay him his wage w. Otherwise, she must pay him w. In the nominal version of the game, $\tilde{v}$, w, $\tilde{g}$, and $\tilde{h}$ are commonly known with certainty by the players. In practice, it seems reasonable that the opportunity costs and the value of work output (e.g., subject to unpredictable defects) would be subject to uncertainty. To this end, suppose that $\tilde{v}$, $\tilde{g}$, and $\tilde{h}$ are subject to independent uncertainty, the nature of which is common knowledge between the two players. For example, we may consider the robust game in which the payoff uncertainty set is given by

$$U = \left\{ \begin{pmatrix} (0, -\tilde{h}) & (w, -w) \\ (w - \tilde{g}, \tilde{v} - w - \tilde{h}) & (w - \tilde{g}, \tilde{v} - w) \end{pmatrix} \;\middle|\; \left(\tilde{g}, \tilde{v}, \tilde{h}\right) \in [\underline{g}, \overline{g}] \times [\underline{v}, \overline{v}] \times [\underline{h}, \overline{h}] \right\}.$$

Example 2. Robust Free-Rider Problem

Consider the symmetric version of the classical, 2-player, free-rider problem discussed in [54]. Each player must make a binary decision of whether or not (actions 1 and 2, respectively) to contribute to the construction of a public good. The players make their decisions simultaneously. If a player contributes, the player incurs some cost $\tilde{c}$, which is subject to minor uncertainty (e.g., because projected costs are rarely accurate), in a way that is common knowledge to the two players. If the public good is built, each player enjoys a payoff of 1. The good will not be built unless at least one player contributes. So, we may consider the resulting robust game with payoff uncertainty set

$$U = \left\{ \left(\begin{array}{cc} (1-\tilde{c}, 1-\tilde{c}) & (1-\tilde{c}, 1) \\ (1, 1-\tilde{c}) & (0,0) \end{array} \right) \middle| \tilde{c} \in [\check{c} - \Delta, \check{c} + \Delta] \right\},$$

for some fixed $\Delta > 0$.

Example 3. Robust Network Routing

Network routing games, formulated as early as 1952 by Wardrop [164], have become an increasingly popular topic in the game theory literature. Within the last five years, Papadimitriou [120] and others have studied the so-called "price of anarchy," or the difference between total payoffs at equilibria versus at Pareto optimality.

Consider a network routing game, in which N internet service providers must each contract for the use of a single "path" in a network of a paths (e.g., servers, wiring, etc.). The providers must make these arrangements simultaneously and prior to knowing the demand to be faced (i.e., the amount of data their customers will want to route). So, each provider's action space is the set of paths in the network. Suppose edge latencies in the network are linear and additive, and that the payoff to provider i when he uses path j_i is given by the negative of total latency experienced on edge j_i. That is, higher latencies yield lower payoffs. Specifically, we can express the uncertain payoff matrix $\tilde{\mathbf{P}}$ as a function, $\mathbf{P}$, of the uncertain demands to be faced. Let $\tilde{d}_i$ denote the uncertain demand to be faced by provider i. $\forall i \in \{1, \ldots, N\}$, $\forall (j_1, \ldots, j_N) \in \{1, \ldots, a\}^N$, let

$$P^i_{(j_1,\ldots,j_N)}\left(\tilde{d}_1, \ldots, \tilde{d}_N\right) = - \sum_{\{i' \mid j_{i'}=j_i\}} \lambda_{(i',j_{i'})} \tilde{d}_{i'},$$

where $\lambda_{(i,j_i)}$ are nonnegative coefficients that account for the fact that the marginal latencies may differ by provider and path. The demand uncertainty may arise from the fact that the providers commonly know the total demand D to be faced by all of them, but do not know how this demand will be distributed among them (e.g., uncertainty of projected subscribership for a future year). For example, the uncertainty set may be given by

$$U = \left\{ \mathbf{P}\left(\tilde{d}_1, \ldots, \tilde{d}_N\right) \;\middle|\; \sum_{i=1}^{N} \tilde{d}_i = D,\ \tilde{d}_i \geq \underline{d_i},\ i = 1, \ldots, N \right\},$$

where

$$\begin{aligned} D &> \sum_{i=1}^{N} \underline{d_i} \\ \underline{d_i} &> 0, \qquad\qquad\qquad i = 1, \ldots, N, \end{aligned}$$

are commonly known by the players.

4.2.6 Nonexistence of *Ex Post* Equilibria

We will now use the incomplete-information inspection game presented in Example 1 to illustrate that not all incomplete-information games have an *ex post* equilibrium. Each possible realization of $\left(\tilde{g}, \tilde{v}, \tilde{h}\right) \in [\underline{g}, \overline{g}] \times [\underline{v}, \overline{v}] \times [\underline{h}, \overline{h}]$ gives rise to a nominal game, i.e., a game with complete information. It is easy to show that each such game has a unique equilibrium in which the employee shirks (action 1) with probability $\tilde{h}/w$ and the employer inspects (action 1) with probability $\tilde{g}/w$. So, unless $\underline{g} = \overline{g}$ and $\underline{h} = \overline{h}$, this incomplete-information game has no *ex post* equilibria.

Accordingly, the *ex post* equilibrium concept cannot be applied to all games with incomplete information, because such equilibria need not exist. In contrast, in the next section, we will prove that any robust finite game with bounded uncertainty set has an equilibrium. In this way, robust games offer an alternative distribution-free notion of equilibrium, whose existence is guaranteed.

4.3 Existence of Equilibria in Robust Finite Games

Having formalized and given examples illustrating our robust optimization model of games with incomplete payoff information, let us now establish the existence of equilibria in the resulting robust games, when these games are finite and have bounded uncertainty sets. Our proof of existence directly uses Kakutani's Fixed Point Theorem [79] and parallels Nash's first existence proof in [116]. As already mentioned, we focus in this section on incomplete-information games not involving private information. In Section 4.6, we extend our existence result to the general case involving potentially private information.

To begin, let us state Kakutani's theorem and a relevant definition. Kakutani's definition of upper semi-continuity relates to mappings from a closed, bounded, convex set S in a Euclidean space into the family of all closed, convex subsets of S. 2^S will denote the power set of S.

Definition 4.3.1 (Kakutani [79]). A point-to-set mapping $\Psi : S \rightarrow 2^S$ is called **upper semi-continuous** if

$$\begin{aligned} \mathbf{y}^n \in & \ \Psi(\mathbf{x}^n), & n = 1, 2, 3, \ldots \\ \lim_{n\rightarrow\infty} \mathbf{x}^n = & \ \mathbf{x} & \\ \lim_{n\rightarrow\infty} \mathbf{y}^n = & \ \mathbf{y} & \end{aligned}$$

imply that $\mathbf{y} \in \Psi(\mathbf{x})$. In other words, the graph of $\Psi(\mathbf{x})$ must be closed.

Theorem 4.3.1 (Kakutani's Fixed Point Theorem [79]). *If S is a closed, bounded, and convex set in a Euclidean space, and Φ is an upper semi-continuous point-to-set mapping of S into the family of closed, convex subsets of S, then $\exists \mathbf{x} \in S$ s.t. $\mathbf{x} \in \Phi(\mathbf{x})$.*

To use Kakutani's Fixed Point Theorem, we must first establish some properties of the worst-case expected payoff functions, given by

$$\rho_i\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right) \triangleq \inf_{\tilde{\mathbf{P}} \in U} \pi_i\left(\tilde{\mathbf{P}}; \mathbf{x}^1, \ldots, \mathbf{x}^N\right), \tag{4.7}$$

$i \in \{1, \ldots, N\}$. In an N-person, robust finite game, let $U \subseteq \mathbb{R}^{N \prod_{i=1}^{N} a_i}$ be the uncertainty set of possible payoff matrices $\tilde{\mathbf{P}}$.

Lemma 4.3.1. *Let $U \subseteq \mathbb{R}^{N \prod_{i=1}^{N} a_i}$ be bounded. Then, $\forall \left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right) \in \mathbb{R}^{a_1 + \cdots + a_N}$ and $\forall \epsilon > 0$, $\exists \delta\left(\epsilon, \mathbf{x}^1, \ldots, \mathbf{x}^N\right) > 0$ such that, $\forall \tilde{\mathbf{P}} \in U$ and $\forall i \in \{1, \ldots, N\}$,*

$$\left\| \left(\mathbf{y}^1, \ldots, \mathbf{y}^N\right) - \left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right) \right\|_\infty < \delta\left(\epsilon, \mathbf{x}^1, \ldots, \mathbf{x}^N\right)$$

implies

$$\left| \pi_i \left(\tilde{\mathbf{P}}; \mathbf{y}^1, \ldots, \mathbf{y}^N\right) - \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^1, \ldots, \mathbf{x}^N\right) \right| < \epsilon.$$

Proof. $\forall \left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right) \in \mathbb{R}^{a_1 + \cdots + a_N}$, and $\forall \epsilon > 0$, consider $\delta\left(\epsilon, \mathbf{x}^1, \ldots, \mathbf{x}^N\right)$ given by

$$\delta\left(\epsilon, \mathbf{x}^1, \ldots, \mathbf{x}^N\right) = \frac{\min\{\epsilon, 1\}}{2\left(2^N - 1\right) M \cdot \prod_{i=1}^{N} \left(a_i \max\left\{\max_{j_i \in \{1, \ldots, a_i\}} \left|x^i_{j_i}\right|, 1\right\}\right)},$$

where $1 < M < \infty$ satisfies

$$\left|\tilde{P}^i_{(j_1, \ldots, j_N)}\right| \leq M, \qquad \forall i \in \{1, \ldots, N\},\ \forall (j_1, \ldots, j_N) \in \prod_{i=1}^{N} \{1, \ldots, a_i\},\ \forall \tilde{\mathbf{P}} \in U.$$

The result follows from algebraic manipulation. □

Lemma 4.3.1 immediately gives rise to the following continuity result, which we therefore state without proof.

Lemma 4.3.2. *Let $U \subseteq \mathbb{R}^{N \prod_{i=1}^{N} a_i}$ be bounded. Then, $\forall i \in \{1, \ldots, N\}$, $\rho_i\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right)$ is continuous on $\mathbb{R}^{a_1 + \cdots + a_N}$.*

Similarly, it is trivial to prove the following lemma.

Lemma 4.3.3. *$\forall i \in \{1, \ldots, N\}$ and $\forall \mathbf{x}^{-i} \in S_{-i}$ fixed, $\rho_i\left(\mathbf{x}^{-i}, \mathbf{x}^i\right)$ is concave in $\mathbf{x}^i$.*

We are now ready to prove the existence of equilibria in robust finite games with bounded uncertainty sets.

Theorem 4.3.2 (Existence of Equilibria in Robust Finite Games). *Any N-person, non-cooperative, simultaneous-move, one-shot robust game, in which $N < \infty$, in which player $i \in \{1, \dots, N\}$ has $1 < a_i < \infty$ possible actions, in which the uncertainty set of payoff matrices $U \subseteq \mathbb{R}^{N \prod_{i=1}^{N} a_i}$ is bounded, and in which there is no private information, has an equilibrium.*

Proof. We will proceed by constructing a point-to-set mapping that satisfies the conditions of Kakutani's Fixed Point Theorem [79], and whose fixed points are precisely the equilibria of the robust game. To begin, clearly, S is closed, bounded, and convex, since S_{a_i} is, $\forall i \in \{1, \dots, N\}$. Define $\Phi : S \to 2^S$ as

$$\begin{aligned} &\Phi\left(\mathbf{x}^1, \dots, \mathbf{x}^N\right) \\ &\quad = \left\{ \left(\mathbf{y}^1, \dots, \mathbf{y}^N\right) \in S \;\middle|\; \mathbf{y}^i \in \arg\max_{\mathbf{u}^i \in S_{a_i}} \rho_i\left(\mathbf{x}^{-i}, \mathbf{u}^i\right),\ i = 1, \dots, N \right\}. \end{aligned} \tag{4.8}$$

Let us show that $\Phi\left(\mathbf{x}^1, \dots, \mathbf{x}^N\right) \neq \emptyset$, $\forall \left(\mathbf{x}^1, \dots, \mathbf{x}^N\right) \in S$. By Lemma 4.3.2, $\forall i \in \{1, \dots, N\}$, $\forall \mathbf{x}^{-i} \in S_{-i}$ fixed, $\rho_i\left(\mathbf{x}^1, \dots, \mathbf{x}^N\right)$ is continuous on S_{a_i}, a nonempty, closed, and bounded subset of $\mathbb{R}^{a_i}$. Thus, by Weierstrass' Theorem,

$$\arg\max_{\mathbf{u}^i \in S_{a_i}} \rho_i\left(\mathbf{x}^{-i}, \mathbf{u}^i\right) \neq \emptyset.$$

Accordingly, $\forall \left(\mathbf{x}^1, \dots, \mathbf{x}^N\right) \in S$,

$$\Phi\left(\mathbf{x}^1, \dots, \mathbf{x}^N\right) \neq \emptyset.$$

It is obvious from the definition of Φ, that $\forall \left(\mathbf{x}^1, \dots, \mathbf{x}^N\right) \in S$, $\Phi\left(\mathbf{x}^1, \dots, \mathbf{x}^N\right) \subseteq S$, and that $\left(\mathbf{x}^1, \dots, \mathbf{x}^N\right)$ is an equilibrium of the robust game iff it is a fixed point of Φ. Thus, we need only prove the existence of a fixed point of Φ. Let us therefore establish that Φ satisfies the remaining conditions of Kakutani's Fixed Point Theorem; that is, we must show that Φ maps S into a family of closed, convex sets, and that Φ is upper semi-continuous.

Let us first prove that, $\forall \left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right) \in S$, $\Phi\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right)$ is a convex set. Suppose

$$\left(\mathbf{u}^1, \ldots, \mathbf{u}^N\right), \left(\mathbf{v}^1, \ldots, \mathbf{v}^N\right) \;\; \in \;\; \Phi\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right).$$

Then, by the definition of Φ, $\forall i \in \{1, \ldots, N\}$, $\forall \mathbf{y}^i \in S_{a_i}$,

$$\rho_i\left(\mathbf{x}^{-i}, \mathbf{u}^i\right) \;\; = \;\; \rho_i\left(\mathbf{x}^{-i}, \mathbf{v}^i\right) \;\; \geq \;\; \rho_i\left(\mathbf{x}^{-i}, \mathbf{y}^i\right).$$

It follows that, $\forall \lambda \in [0, 1]$, $\forall \mathbf{y}^i \in S_{a_i}$,

$$\lambda \rho_i\left(\mathbf{x}^{-i}, \mathbf{u}^i\right) + (1-\lambda)\rho_i\left(\mathbf{x}^{-i}, \mathbf{v}^i\right) \;\; \geq \;\; \rho_i\left(\mathbf{x}^{-i}, \mathbf{y}^i\right).$$

By the concavity result of Lemma 4.3.3,

$$\lambda(\mathbf{u}^1, \ldots, \mathbf{u}^N) + (1-\lambda)(\mathbf{v}^1, \ldots, \mathbf{v}^N) \;\; \in \;\; \Phi\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right).$$

Let us now show that Φ is upper semi-continuous, per Kakutani's definition. Suppose that, for $n = 1, 2, 3, \ldots$,

$$\begin{aligned}
\left(\mathbf{x}^{1,n}, \ldots, \mathbf{x}^{N,n}\right) &\;\; \in \;\; S \\
\left(\mathbf{y}^{1,n}, \ldots, \mathbf{y}^{N,n}\right) &\;\; \in \;\; \Phi\left(\mathbf{x}^{1,n}, \ldots, \mathbf{x}^{N,n}\right) \\
\lim_{n\to\infty}\left(\mathbf{x}^{1,n}, \ldots, \mathbf{x}^{N,n}\right) &\;\; = \;\; \left(\mathbf{u}^1, \ldots, \mathbf{u}^N\right) \;\; \in \;\; S \\
\lim_{n\to\infty}\left(\mathbf{y}^{1,n}, \ldots, \mathbf{y}^{N,n}\right) &\;\; = \;\; \left(\mathbf{v}^1, \ldots, \mathbf{v}^N\right) \;\; \in \;\; S.
\end{aligned}$$

By the definition of Φ, we know that, $\forall n = 1, 2, 3, \ldots$, $\forall i \in \{1, \ldots, N\}$ and $\forall \mathbf{w}^i \in S_{a_i}$,

$$\rho_i\left(\mathbf{x}^{-i,n}, \mathbf{y}^{i,n}\right) \;\; \geq \;\; \rho_i\left(\mathbf{x}^{-i,n}, \mathbf{w}^i\right).$$

Taking the limit of both sides, and using Lemma 4.3.2 (continuity of ρ_i), we obtain that, $\forall i \in \{1, \dots, N\}$ and $\forall \mathbf{w}^i \in S_{a_i}$,

$$\rho_i \left(\mathbf{u}^{-i}, \mathbf{v}^i\right) \quad \geq \quad \rho_i \left(\mathbf{u}^{-i}, \mathbf{w}^i\right).$$

Hence,

$$\left(\mathbf{v}^1, \dots, \mathbf{v}^N\right) \quad \in \quad \Phi\left(\mathbf{u}^1, \dots, \mathbf{u}^N\right),$$

and Φ is upper semi-continuous. Note that the fact that $\Phi\left(\mathbf{x}^1, \dots, \mathbf{x}^N\right)$ is closed follows from the fact that Φ is upper semi-continuous (take $\left(\mathbf{x}^{1,n}, \dots, \mathbf{x}^{N,n}\right) = \left(\mathbf{u}^1, \dots, \mathbf{u}^N\right) = \left(\mathbf{x}^1, \dots, \mathbf{x}^N\right)$, $\forall n = 1, 2, 3, \dots$).

This completes the proof that Φ satisfies the conditions of Kakutani's Fixed Point Theorem, and thereby establishes the existence of an equilibrium in the robust game. □

4.4 Computing Sample Equilibria of Robust Games

In Section 4.3, we established the existence of robust-optimization equilibria in any finite, incomplete-information game with bounded uncertainty set and no private information. In this section, for any resulting robust game with bounded polyhedral uncertainty set, we show that the set of equilibria is a projection of the set of solutions to a system of multilinear equalities and inequalities. This projection is a simple component-wise one, into a space of lower dimension. Based on this formulation, we present an approximate computation method for finding a sample equilibrium of such a robust game. We provide numerical results from the application of our method. Finally, we describe a class of robust finite games whose set of equilibria are precisely the set of equilibria in a related complete-information, finite game in the same action spaces. As noted before, in this section, we focus on robust games not involving private information. In Section 4.6, we provide a more general result on the computation of robust-optimization equilibria in games with private information.

4.4.1 Review for Complete-Information, Finite Games

Before describing our technique for finding robust-optimization equilibria, let us review the state of the art for complete-information, finite games.

Solving for an equilibrium of a general complete-information, finite game is regarded as a difficult task [120]. Two-person, zero-sum games are the exception. As noted in von Stengel [162], in any such game, the set of Nash equilibria is precisely the set of maximinimizers, as defined by von Neumann and Morgenstern [161]. Accordingly, the equilibria are pairs of solutions of two separate LPs, one for each player, and the set of equilibria is therefore convex. For non-fixed-sum games, solving for Nash equilibria is more computationally demanding, and the set of equilibria is generally nonconvex. As discussed in McKelvey and McLennan [107], the set of Nash equilibria can be cast as the solution set of several well-known problems in the optimization literature: fixed point problems, nonlinear complementarity problems (linear in the case of two-player games), stationary point problems, systems of multilinear equalities and inequalities, and unconstrained penalty function minimization problems, in which a penalty is incurred for violations of the multilinear constraints.

Algorithms for finding sample Nash equilibria exploit special properties of these formulations. Traditionally, the most well-regarded algorithm for two-player, non-fixed-sum, complete-information, finite games has been the Lemke-Howson path-following algorithm [90] for linear complementarity problems. For this more general class of problems, the algorithm's worst-case runtime is exponential. The worst-case runtime in the specific application context of two-person games is unknown. For N-player complete-information, finite games with $N > 2$, the traditionally favored algorithms have been different versions of path-following methods based on Scarf's simplicial subdivision approach [145, 146] to computing fixed points of a continuous function on a compact set. These simplicial subdivision algorithms include that of van der Laan and Talman [158, 159], and their worst-case runtimes are also exponential. The Lemke-Howson and simplicial subdivision algorithms form the backbone of the well-known game theory software Gambit [108].

More recent approaches to solving for sample Nash equilibria have exploited the multilinear system formulation, and have applied general root-finding methods to the complementarity conditions that arise from this system. For an overview of the formulation, we refer the reader to Chapter 6 of [153]. For a comparison of these Gröbner basis and homotopy continuation methods of computation with the more traditional Gambit software, we refer the reader to Datta [40]. Govindan's and Wilson's [60] global Newton method similarly uses the multilinear system formulation. Finally, Porter, Nudelman, and Shoham [130] offer a potential shortcut, which exploits the fact that, for complete-information games, it is easier to solve for a Nash equilibrium with a fixed support, and that smaller supports yield lower runtimes.

These more recent numerical techniques are more powerful in terms of their aptitudes at computing all equilibria of a complete-information, finite game, a more difficult task than computing a single, sample equilibrium. For example, PHCpack [160] can find all isolated roots of a system of polynomials.

4.4.2 Robust Finite Games

Multilinear System Formulation for Equilibria

In this subsection, we show that the set of equilibria of a robust finite game, with bounded polyhedral uncertainty set and no private information, is the projection of the solution set of a system of multilinear equalities and inequalities. The projection is a simple component-wise projection into a space of lower dimension.

As a basis of comparison, for an N-player, complete-information, finite game, in which player i has action space $\{1, \ldots, a_i\}$, and in which the payoff matrix is $\check{\mathbf{P}}$, let us formulate the multilinear system whose solutions are the set of Nash equilibria. From Condition (4.3), we see that $\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right)$ is a Nash equilibrium iff it satisfies the following system:

$$
\begin{aligned}
\pi_i\left(\check{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i - \mathbf{e}^i_{j_i}\right) &\geq 0, && i = 1, \ldots, N;\ j_i = 1, \ldots, a_i \\
\mathbf{e}'\mathbf{x}^i &= 1, && i = 1, \ldots, N \\
\mathbf{x}^i &\geq \mathbf{0}, && i = 1, \ldots, N,
\end{aligned}
$$

where $\mathbf{e}$ is the vector, of appropriate dimension, of all ones, and where $\mathbf{e}^i_{j_i}$ denotes the j_i^{th} unit vector in $\mathbb{R}^{a_i}$.

Analogously, from Condition (4.5), $\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right)$ is an equilibrium of the robust finite game with closed and bounded uncertainty set $U \subseteq \mathbb{R}^{N \prod_{i=1}^N a_i}$, and with no private information, iff

$$
\begin{aligned}
\min_{\tilde{\mathbf{P}} \in U} \pi_i\left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i\right) - \max_{\mathbf{u}^i \in S_{a_i}} \min_{\tilde{\mathbf{P}} \in U} \pi_i\left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right) &\geq 0, && i = 1, \ldots, N \\
\mathbf{e}'\mathbf{x}^i &= 1, && i = 1, \ldots, N \\
\mathbf{x}^i &\geq \mathbf{0}, && i = 1, \ldots, N.
\end{aligned}
$$

Stated in another way, $\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right)$ is an equilibrium of the robust finite game iff, for each $i \in \{1, \ldots, N\}$, $\mathbf{x}^i$ is a max-min strategy in a two-person, zero-sum game between player i and an adversary. In this two-person, zero-sum game, the payoff matrix is determined by $\mathbf{x}^{-i}$, and the adversary's strategy space is U.

Although the above system may not seem amenable to reformulation as a system of multilinear equalities and inequalities, we establish in Theorem 4.4.1 that, when U is a bounded polyhedron, the system can, in fact, be reformulated in this way. Before stating and proving this theorem, let us state and prove the following lemma, inspired by the LP duality proof techniques used in [18].

Lemma 4.4.1. *Let $U \subseteq \mathbb{R}^{N \prod_{i=1}^N a_i}$ be a bounded polyhedral set, given by*

$$
U = \left\{ \tilde{\mathbf{P}} \;\middle|\; \mathbf{F} \cdot \operatorname{vec}\left(\tilde{\mathbf{P}}\right) \geq \mathbf{d} \right\} \neq \emptyset, \tag{4.9}
$$

where

$$\text{vec}(\mathbf{P}) \triangleq \left(P^i_{(j_1,\ldots,j_N)}\right)_{i=1,\ldots,N;\ (j_1,\ldots,j_N)\in\prod_{i=1}^N\{1,\ldots,a_i\}}.$$

Let $\mathbf{G}(\ell)$, $\ell \in \{1,\ldots,k\}$, *denote the extreme points of* U. $\forall i \in \{1,\ldots,N\}$, $\forall\left(\mathbf{x}^{-i},\mathbf{u}^i\right) \in S$, *the following three conditions are equivalent.*

Condition 1) $z_i \le \min_{\tilde{\mathbf{P}}\in U} \pi_i\left(\tilde{\mathbf{P}};\mathbf{x}^{-i},\mathbf{u}^i\right)$

Condition 2) $z_i \le \pi_i\left(\mathbf{G}(\ell);\mathbf{x}^{-i},\mathbf{u}^i\right)$, $\ell = 1,\ldots,k$.

Condition 3) $\exists \boldsymbol{\eta}^i \in \mathbb{R}^m$ *such that*

$$\begin{aligned} \mathbf{d}'\boldsymbol{\eta}^i &\ge z_i \\ \mathbf{F}'\boldsymbol{\eta}^i &= \mathbf{Y}^i\left(\mathbf{x}^{-i}\right)\mathbf{u}^i \\ \boldsymbol{\eta}^i &\ge \mathbf{0}, \end{aligned}$$

where $\mathbf{Y}^i\left(\mathbf{x}^{-i}\right) \in \mathbb{R}^{\left(N\prod_{i=1}^N a_i\right)\times a_i}$ *denotes the matrix such that*

$$\text{vec}(\mathbf{P})'\,\mathbf{Y}^i\left(\mathbf{x}^{-i}\right)\mathbf{u}^i = \pi_i\left(\mathbf{P};\mathbf{x}^{-i},\mathbf{u}^i\right). \tag{4.10}$$

Proof. Conditions 1 and 2 are equivalent, since by the linearity of π_i in $\tilde{\mathbf{P}}$,

$$\min_{\tilde{\mathbf{P}}\in U} \pi_i\left(\tilde{\mathbf{P}};\mathbf{x}^{-i},\mathbf{u}^i\right) = \min_{\ell\in\{1,\ldots,k\}} \pi_i\left(\mathbf{G}(\ell);\mathbf{x}^{-i},\mathbf{u}^i\right).$$

To prove the equivalence of Conditions 1 and 3, consider the following primal-dual pair, in which $\left(\mathbf{x}^{-i},\mathbf{u}^i\right)$ is treated as data.

$$\begin{aligned} \min_{\text{vec}(\mathbf{P})} \quad & \text{vec}(\mathbf{P})'\,\mathbf{Y}^i\left(\mathbf{x}^{-i}\right)\mathbf{u}^i \\ \text{s.t.} \quad & \mathbf{F}\cdot\text{vec}(\mathbf{P}) \ge \mathbf{d} \end{aligned} \tag{4.11}$$

$$
\begin{aligned}
\max_{\boldsymbol{\eta}^i} \quad & \mathbf{d}'\boldsymbol{\eta}^i \\
\text{s.t.} \quad & \mathbf{F}'\boldsymbol{\eta}^i = \mathbf{Y}^i\left(\mathbf{x}^{-i}\right)\mathbf{u}^i \\
& \boldsymbol{\eta}^i \geq \mathbf{0}.
\end{aligned} \tag{4.12}
$$

Since $U \neq \emptyset$, Problem (4.11) is feasible. Suppose $(\mathbf{x}^{-i}, \mathbf{u}^i)$ satisfies Condition 1. Then, Problem (4.11) is also bounded. By strong duality, Problem (4.12) is feasible and bounded with optimal value equal to that of Problem (4.11). Thus, Condition 3 is satisfied. For the other direction, suppose Condition 3 is satisfied. Then, Problem (4.12) is feasible. By weak duality, Condition 1 must hold. □

Theorem 4.4.1 (Computation of Equilibria in Robust Finite Games). *Consider the N-player robust game, in which player $i \in \{1, \ldots, N\}$ has action set $\{1, \ldots, a_i\}$, $1 < a_i < \infty$, in which the payoff uncertainty set $U \subseteq \mathbb{R}^{N\prod_{i=1}^{N} a_i}$ is polyhedral, bounded, and given by (4.9), and in which there is no private information. Let $\mathbf{G}(\ell)$, $\ell \in \{1, \ldots, k\}$, denote the extreme points of U. The following three conditions are equivalent.*

Condition 1) *$\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right)$ is an equilibrium of the robust game.*

Condition 2) *$\forall i \in \{1, \ldots, N\}$, $\exists z_i \in \mathbb{R}, \boldsymbol{\theta}^i \in \mathbb{R}^k, \phi_i \in \mathbb{R}$ such that $\left(\mathbf{x}^1, \ldots, \mathbf{x}^N, z_i, \boldsymbol{\theta}^i, \phi_i\right)$ satisfies*

$$
\begin{aligned}
z_i &= \phi_i \\
z_i - \pi_i\left(\mathbf{G}(\ell); \mathbf{x}^1, \ldots, \mathbf{x}^N\right) &\leq 0, && \ell = 1, \ldots, k \\
\mathbf{e}'\mathbf{x}^i &= 1 \\
\mathbf{x}^i &\geq \mathbf{0} \\
\mathbf{e}'\boldsymbol{\theta}^i &= 1 \\
\sum_{\ell=1}^{k} \theta_\ell^i \pi_i\left(\mathbf{G}(\ell); \mathbf{x}^{-i}, \mathbf{e}_{j_i}^i\right) - \phi_i &\leq 0, && j_i = 1, \ldots, a_i \\
\boldsymbol{\theta}^i &\geq \mathbf{0},
\end{aligned} \tag{4.13}
$$

where $\mathbf{e}$ *is the vector, of appropriate dimension, of all ones, and where* $\mathbf{e}^i_{j_i}$ *is the* j_i^{th} *unit vector in* $\mathbb{R}^{a_i}$.

Condition 3) $\forall i \in \{1, \dots, N\}$, $\exists \boldsymbol{\eta}^i \in \mathbb{R}^m, \boldsymbol{\xi}^i \in \mathbb{R}^{N \prod_{i=1}^N a_i}$ *such that* $\left(\mathbf{x}^1, \dots, \mathbf{x}^N, \boldsymbol{\eta}^i, \boldsymbol{\xi}^i\right)$ *satisfies*

$$\begin{aligned} \left(\boldsymbol{\xi}^i\right)' \mathbf{Y}^i \left(\mathbf{x}^{-i}\right) \mathbf{e}^i_{j_i} &\leq \mathbf{d}' \boldsymbol{\eta}^i, \qquad && j_i = 1, \dots, a_i \\ \mathbf{F}' \boldsymbol{\eta}^i - \mathbf{Y}^i \left(\mathbf{x}^{-i}\right) \mathbf{x}^i &= \mathbf{0} && \\ \mathbf{e}' \mathbf{x}^i &= 1 && \\ \mathbf{x}^i &\geq \mathbf{0} && \\ \boldsymbol{\eta}^i &\geq \mathbf{0} && \\ \mathbf{F} \boldsymbol{\xi}^i &\geq \mathbf{d}, && \end{aligned} \tag{4.14}$$

where $\mathbf{Y}^i\left(\mathbf{x}^{-i}\right) \in \mathbb{R}^{\left(N \prod_{i=1}^N a_i\right) \times a_i}$ *is as defined in (4.10).*

Proof. Since U is closed and bounded, Condition 1 is equivalent, by Relation (4.5), to

$$\begin{aligned} \left(\mathbf{x}^1, \dots, \mathbf{x}^N\right) &\in S \\ \mathbf{x}^i &\in \arg\max_{\mathbf{u}^i \in S_{a_i}} \left[\min_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right) \right], \qquad i = 1, \dots, N. \end{aligned}$$

In turn, these constraints are equivalent to the requirement that, $\forall i \in \{1, \dots, N\}$, $\exists z_i \in \mathbb{R}$ such that $(\mathbf{x}^i, z_i)$ is a maximizer of the following robust LP.

$$\begin{aligned} \max_{\mathbf{u}^i, z_i} \quad & z_i \\ \text{s.t.} \quad & z_i \leq \min_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right) \\ & \mathbf{e}' \mathbf{u}^i = 1 \\ & \mathbf{u}^i \geq \mathbf{0}. \end{aligned} \tag{4.15}$$

In this robust LP, $\mathbf{x}^{-i}$ is regarded as given data, and $\mathbf{e}$ denotes the vector, of appropriate dimension, of all ones.

Suppose Condition 1 is satisfied. Then, by Lemma 4.4.1, $\exists z_i \in \mathbb{R}$ and $\boldsymbol{\eta}^i \in \mathbb{R}^m$ such that $(\mathbf{x}^i, z_i)$ is an optimal solution of the maximization problem in the primal-dual pair

$$\begin{aligned} \max_{\mathbf{u}^i, z_i} \; & z_i \\ \text{s.t. } z_i \; & \leq \; \pi_i\left(\mathbf{G}(\ell); \mathbf{x}^{-i}, \mathbf{u}^i\right), \qquad \ell = 1, \ldots, k \\ \mathbf{e}'\mathbf{u}^i \; & = \; 1 \\ \mathbf{u}^i \; & \geq \; \mathbf{0}, \end{aligned} \tag{4.16}$$

$$\begin{aligned} \min_{\boldsymbol{\theta}^i, \phi_i} \; & \phi_i \\ \text{s.t. } \mathbf{e}'\boldsymbol{\theta}^i \; & = \; 1 \\ \sum_{\ell=1}^{k} \theta^i_\ell \pi_i\left(\mathbf{G}(\ell); \mathbf{x}^{-i}, \mathbf{e}^i_{j_i}\right) - \phi_i \; & \leq \; 0, \qquad j_i = 1, \ldots, a_i \\ \boldsymbol{\theta}^i \; & \geq \; \mathbf{0}, \end{aligned} \tag{4.17}$$

and such that $(\mathbf{x}^i, \boldsymbol{\eta}^i, z_i)$ is an optimal solution of the maximization problem in the primal-dual pair

$$\begin{aligned} \max_{\mathbf{u}^i, \boldsymbol{\eta}^i, z_i} \; & z_i \\ \text{s.t. } z_i - \mathbf{d}'\boldsymbol{\eta}^i \; & \leq \; 0 \\ \mathbf{F}'\boldsymbol{\eta}^i - \mathbf{Y}^i\left(\mathbf{x}^{-i}\right)\mathbf{u}^i \; & = \; 0 \\ \mathbf{e}'\mathbf{u}^i \; & = \; 1 \\ \mathbf{u}^i, \; \boldsymbol{\eta}^i \; & \geq \; \mathbf{0}, \end{aligned} \tag{4.18}$$

$$\begin{aligned} \min_{\boldsymbol{\xi}^i, \nu_i} \; & \nu_i \\ \text{s.t. } \mathbf{F}\boldsymbol{\xi}^i \; & \geq \; \mathbf{d} \\ \nu_i \; & \geq \; \left(\boldsymbol{\xi}^i\right)' \mathbf{Y}^i\left(\mathbf{x}^{-i}\right)\mathbf{e}^i_{j_i}, \qquad j_i = 1, \ldots, a_i. \end{aligned} \tag{4.19}$$

	robust game using constraints	robust game using extreme points	complete-info game
variables	$a_{\text{tot}} + N(m+v)$	$a_{\text{tot}} + N(k+2)$	a_{tot}
constraints	$2a_{\text{tot}} + N(2m+v+1)$	$2a_{\text{tot}} + N(2k+3)$	$2a_{\text{tot}} + N$
maximum degree	N	N	N

Table 4.1: Sizes of multilinear systems for equilibria

Conditions 2 and 3 follow from LP strong duality.

For the reverse direction, suppose that Condition 2 holds. Then, $\forall i \in \{1,\ldots,N\}$, and for $\mathbf{x}^{-i}$, $(\mathbf{x}^i, z_i)$ is a feasible solution of (4.16), and $(\boldsymbol{\theta}^i, \phi_i)$ is a feasible solution of (4.17), such that $z_i = \phi_i$. By LP weak duality, $(\mathbf{x}^i, z_i)$ is an optimizer of (4.16). Equivalently, by Lemma 4.4.1, $(\mathbf{x}^i, z_i)$ is an optimizer of (4.15), and Condition 1 follows.

Similarly, suppose that Condition 3 holds. $\forall i \in \{1,\ldots,N\}$, let

$$\begin{aligned} z_i &= \mathbf{d}'\boldsymbol{\eta}^i \\ \nu_i &= \max_{j_i \in \{1,\ldots,a_i\}} \left(\boldsymbol{\xi}^i\right)' \mathbf{Y}^i\left(\mathbf{x}^{-i}\right)\mathbf{e}^i_{j_i}. \end{aligned}$$

Then, for $\mathbf{x}^{-i}$, $(\mathbf{x}^i, \boldsymbol{\eta}^i, z_i)$ is a feasible solution of (4.18) and $(\boldsymbol{\xi}^i, \nu_i)$ is a feasible solution of (4.19) such that $z_i \geq \nu_i$. By LP weak duality, $(\mathbf{x}^i, \boldsymbol{\eta}^i, z_i)$ is an optimizer of (4.18). Equivalently, by Lemma 4.4.1, $(\mathbf{x}^i, z_i)$ is an optimizer of (4.15), and Condition 1 follows. □

Remark: Note that Systems (4.13) and (4.14) are derived using the extreme-point and constraint representations of the polyhedral set U, respectively. These systems are very sparse as a result of their multilinearity. In addition, it is possible to formulate System (4.14) more compactly if U can be described by m constraints and only v variables, with $v < N\prod_{i=1}^N a_i$. Let $a_{\text{tot}} = \sum_{i=1}^N a_i$, v and m be the number of variables and constraints, respectively, needed to define U, and k be the number of extreme points of U. Table 4.1 summarizes the sizes of the different multilinear systems of equalities and inequalities whose solution sets are precisely the set of equilibria of an N-player game in which player $i \in \{1,\ldots,N\}$ has action space $\{1,\ldots,a_i\}$.

Computation Method

For robust finite games with bounded polyhedral uncertainty sets and no private information, we showed in Section 4.4.2 that the set of equilibria is a projection of the solution set of a system of multilinear equalities and inequalities. This projection is a simple component-wise projection into a space of lower dimension. Currently available and computationally effective solvers for large polynomial systems tend to be specific to systems of *equations* and not inequalities. Accordingly, we propose to solve the multilinear systems for the robust-optimization equilibria by converting any such system into a corresponding penalty function, and then solving the resulting unconstrained minimization problem. The penalty method we use is based on Courant's quadratic loss technique [34], which Fiacco and McCormick later more fully developed in [49].

To more concretely describe our approach, consider any system

$$\begin{aligned} g_n(\mathbf{y}) &= 0, \qquad && n \in E \\ g_n(\mathbf{y}) &\leq 0, \qquad && n \in I, \end{aligned} \tag{4.20}$$

with $\mathbf{y} \in \mathbb{R}^V$, $|I| < \infty$, and $|E| < \infty$. Let

$$h(\mathbf{y}) = \frac{1}{2}\sum_{n\in E}[g_n(\mathbf{y})]^2 + \frac{1}{2}\sum_{n\in I}[\max\{g_n(\mathbf{y}), 0\}]^2 .$$

Since $h(\mathbf{y}) \geq 0$, $\forall \mathbf{y} \in \mathbb{R}^V$, it is easy to see that $\mathbf{y}$ satisfies System (4.20) iff

$$h(\mathbf{y}) = \min_{\mathbf{u}\in\mathbb{R}^V} h(\mathbf{u}) = 0.$$

So, we can solve System (4.20) by solving the unconstrained minimization problem

$$\min_{\mathbf{u}\in\mathbb{R}^V} h(\mathbf{u}).$$

For the unconstrained minimization problem, we propose the use of a pseudo-Newton method using the Armijo rule (see, for example, [15]) for determining step size at each iteration. Each pseudo-Newton method run attempts to find a single point satisfying the constraints. It is possible, though not guaranteed, that, when the system of constraints has more than one solution, multiple pseudo-Newton method runs may identify multiple, distinct approximate solutions of the system. Furthermore, in contrast to most other state-of-the-art polynomial system solvers, this method is capable of finding non-isolated, as well as isolated solutions.

In the next subsection, we present numerical results from the implementation of this technique for approximately computing a sample robust-optimization equilibrium.

Numerical Results

For each problem instance, we formulated the set of equilibria using System (4.13). We executed all computations in MATLAB 6.5.0 R13, running on the Red Hat Linux 7.2-1 operating system, on a Dell with a Pentium IV processor, at 1.7 GHz with 512 MB RAM. To encourage the numerical method to find points satisfying the nonnegativity and normalization constraints on $\mathbf{x}^i$ and $\boldsymbol{\theta}^i$, $i \in \{1, \ldots, N\}$, we multiplied the amount of violation of each such constraint by $M = 100$, before halving the square of this violation. We initialized all runs of the pseudo-Newton method by, for each $i \in \{1, \ldots, N\}$, randomly generating $\mathbf{x}^i$ and $\boldsymbol{\theta}^i$, satisfying the aforementioned nonnegativity and normalization constraints. We initialized z_i to be the maximum possible value satisfying the upper-bound constraint on z_i, and we set ϕ_i either equal to z_i or to the minimum possible value allowed by the lower-bound constraint on ϕ_i. For each pseudo-Newton method run, we terminated the run if the current and previous iterate were too close, if the norm of the direction was too small, if the penalty was already sufficiently small, or if the number of iterations already executed was too large.

We executed the method on the robust inspection game, described in Example 1 in

Section 4.2.5, with

$$\underline{g} = 8, \qquad \underline{v} = 16, \qquad \underline{h} = 4, \qquad w = 15,$$

$$\overline{g} = 12, \qquad \overline{v} = 24, \qquad \overline{h} = 6.$$

The multilinear system for the equilibria of this robust game has 22 constraints in 10 variables, after elimination of some redundant variables. We terminated the pseudo-Newton method run once the penalty function dipped below 10^{-8}. As will follow from Theorem 4.4.2, in the unique equilibrium of this robust game, the employee (row player) shirks (plays action 1) with probability $x_1^1 = 2/5$, and the employer (column player) inspects (plays action 1) with probability $x_1^2 = 4/5$. Our numerical method terminates at $(x_1^1, x_2^2) = (0.4000, 0.8000)$ after 0.5000 seconds of one pseudo-Newton run, requiring 71 iterations.

In addition, we executed the method on the robust free-rider game, described in Example 2 in Section 4.2.5, with

$$\underline{c} = 1/4, \qquad \overline{c} = 5/8.$$

The multilinear system for the equilibria of this robust game has 18 constraints in 8 variables, after elimination of some redundant variables. We terminated each pseudo-Newton method run once the penalty function dipped below 10^{-10} or the number of iterations reached 2000. We used $M = 1$, since the method did not seem to be attracted to strategy profiles outside of the simplex. Let x_1^i denote the probability with which player $i \in \{1, 2\}$ contributes. As will follow from Theorem 4.4.2, this robust game has 3 equilibria (x_1^1, x_1^2): (1,0), (0,1), and $(1 - \overline{c}, 1 - \overline{c}) = (3/8, 3/8)$. We made 15 sequential runs of the pseudo-Newton method, each initialized at a randomly generated point. These 15 runs required 1.8458 minutes, with each run executing an average of 1,652.1 iterations in an average of 7.3827 seconds. Terminal points with penalty function less than 10^{-10} included $(0.0000, 0.9999)$, $(0.9999, 0.0000)$, $(0.3750, 0.3751)$, and $(0.3751, 0.3750)$. This example demonstrates that the method is capable of finding multiple equilibria, and possibly all equilibria, of a robust game.

Lastly, we executed the method on several instances of the robust network routing game, described in Example 3 in Section 4.2.5. The instances differ in terms of their values of N, the number of players, and a, the number of paths available. The resulting versions of System (4.13) consist of $2N^2 + N(2a+3)$ constraints in $N^2 + N(2+a)$ variables.

For all the instances, we used the same values for D and $\boldsymbol{\lambda}$. In particular, we set $D = 5$ and $\lambda_{(i,j_i)}$ to be a realization of the uniform distribution on $[0,4]$. The computational results for these robust network routing games are summarized in Table 4.2. For each instance, we made only one run of the pseudo-Newton method, and terminated it after the lesser of 50 iterations or the minimum number of iterations required to produce an iterate with associated penalty less than 10^{-5}. The "vars" and "constr's" columns in Table 4.2 give the number of variables and constraints, respectively, in System (4.13) for each problem instance. The "iters" column gives the number of iterations executed. The "penalty" column gives the penalty value of the final iterate. Finally, the "proportional error" column gives

$$\frac{\text{penalty}}{\min\limits_{i \in \{1,\ldots,N\}} \min\left\{|\hat{z}_i|, |\hat{\phi}_i|\right\}},$$

where $\hat{z}_i$ and $\hat{\phi}_i$ denote the values of z_i and ϕ_i in the final iterate. We could obviously achieve better speed or accuracy by varying the cap on the number of iterations and the penalty threshold used to decide whether to terminate the pseudo-Newton method run.

These numerical results demonstrate that a practical method, simple in nature and general in its applicability, exists for approximately solving, with considerable accuracy and speed, for sample equilibria of robust games of small size. Furthermore, with longer run-times and lower accuracy, this method may be capable of finding solutions for robust finite games of larger size.

A Special Class of Robust Finite Games

Under certain conditions, the set of equilibria of a robust finite game is equivalent to that of a related finite game with complete payoff information, with the same number of players, and

	vars	constr's	cpu time (mins)	iters	penalty	proportional error
$N = 2,\ a = 2$	12	22	0.0612	37	7.9811×10^{-6}	1.8850×10^{-7}
$N = 3,\ a = 2$	21	39	0.3887	17	5.5489×10^{-8}	8.0825×10^{-10}
$N = 3,\ a = 3$	24	45	1.4572	50	9.7746×10^{-3}	2.2952×10^{-4}
$N = 4,\ a = 2$	32	60	3.4895	50	9.2659×10^{-3}	1.3192×10^{-4}
$N = 4,\ a = 3$	36	68	3.8935	50	1.2910×10^{-1}	3.0427×10^{-3}
$N = 4,\ a = 4$	40	76	4.7893	50	3.8569	1.4566×10^{-1}
$N = 5,\ a = 2$	45	85	7.3268	50	5.7834×10^{-1}	7.7508×10^{-3}
$N = 5,\ a = 3$	50	95	9.2945	50	1.3551	2.5495×10^{-2}
$N = 5,\ a = 4$	55	105	12.6322	50	3.4239	8.7083×10^{-2}
$N = 5,\ a = 5$	60	115	17.6880	50	15.3203	5.9287×10^{-1}

Table 4.2: Numerical results for instances of robust network routing game

with the same action spaces. In these cases, equilibria computation for the robust game will reduce to computation in the context of the related complete-information finite game. As shown in Table 4.1, the multilinear systems arising from robust finite games are larger than those arising from complete-information, finite games with the same number of players and with the same action spaces. Thus, it will be computationally beneficial to take advantage of this equivalence when it holds. As we will discuss in Section 4.5, the complete-information equivalent of the robust finite game will generally not be the nominal (i.e., average) version of the robust game.

The following theorem establishes sufficient conditions for the equivalence of a robust finite game with a complete-information finite game having the same number of players and the same action spaces.

Theorem 4.4.2. *Consider the robust finite game, without private information and in which the payoff uncertainty set is*

$$U = \left\{ \mathbf{P}\left(\tilde{f}_1, \ldots, \tilde{f}_v\right) \;\middle|\; \left(\tilde{f}_1, \ldots, \tilde{f}_v\right) \in U_f \right\}, \tag{4.21}$$

where

$$U_f = \left\{ \left(\tilde{f}_1, \ldots, \tilde{f}_v\right) \middle| \tilde{f}_\ell \in \left[\underline{f_\ell}, \overline{f_\ell}\right], \ \ell \in \{1, \ldots, v\} \right\}, \tag{4.22}$$

and $\mathbf{P}$ *is a continuous and differentiable vector function. Suppose that,* $\forall i \in \{1, \ldots, N\}$ *and* $\forall \ell \in \{1, \ldots, v\}$, $\exists \kappa(i, \ell) \in \{-1, 0, 1\}$ *such that,* $\forall (j_1, \ldots, j_N) \in \prod_{i=1}^{N} \{1, \ldots, a_i\}$ *and* $\forall \left(\tilde{f}_1, \ldots, \tilde{f}_v\right) \in U_f$,

$$sign\left(\frac{\partial}{\partial f_\ell}\left[P^i_{(j_1,\ldots,j_N)}\left(f_1, \ldots, f_v\right)\right]_{(f_1,\ldots,f_v)=\left(\tilde{f}_1,\ldots,\tilde{f}_v\right)}\right) = \kappa(i, \ell).$$

Then $\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right)$ *is an equilibrium of this robust game iff it is a Nash equilibrium of the complete-information, finite game, with the same number of players and the same action spaces, and in which the payoff matrix is* $\mathbf{Q}$, *defined by*

$$\begin{aligned} \mathbf{Q}^i_{(j_1,\ldots,j_N)} &= P^i_{(j_1,\ldots,j_N)}\left(h^i_1, \ldots, h^i_v\right) \\ h^i_\ell &= \begin{cases} \overline{f_\ell}, & \kappa(i,\ell) < 0 \\ \underline{f_\ell}, & \kappa(i,\ell) \geq 0. \end{cases} \end{aligned}$$

Proof. Let $A_i \triangleq \{1, \ldots, a_i\}$ and $A \triangleq \prod_{i=1}^{N} A_i$. $\mathbf{Q} \in U$ implies that, $\forall i \in \{1, \ldots, N\}$, $\forall \left(\mathbf{x}^{-i}, \mathbf{u}^i\right) \in S$, $\pi_i\left(\mathbf{Q}; \mathbf{x}^{-i}, \mathbf{u}^i\right) \geq \min_{\tilde{\mathbf{P}} \in U} \pi_i\left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right)$. Conversely, $\forall i$, $\forall \left(\mathbf{x}^{-i}, \mathbf{u}^i\right) \in S$,

$$\begin{aligned} \min_{\tilde{\mathbf{P}} \in U} \pi_i\left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right) &\geq \sum_{j_1=1}^{a_1} \cdots \sum_{j_i=1}^{a_i} \cdots \sum_{j_N=1}^{a_N} \left(\prod_{\substack{i'=1 \\ i' \neq i}}^{N} x^{i'}_{j_{i'}}\right) u^i_{j_i} \min_{\left(\tilde{f}_1,\ldots,\tilde{f}_v\right) \in U_f} P^i_{(j_1,\ldots,j_N)}\left(\tilde{f}_1, \ldots, \tilde{f}_v\right) \\ &= \sum_{j_1=1}^{a_1} \cdots \sum_{j_i=1}^{a_i} \cdots \sum_{j_N=1}^{a_N} \left(\prod_{\substack{i'=1 \\ i' \neq i}}^{N} x^{i'}_{j_{i'}}\right) u^i_{j_i} P^i_{(j_1,\ldots,j_N)}\left(h^i_1, \ldots, h^i_v\right). \end{aligned}$$

The equality follows since, by the definition of $\mathbf{h}$, $\forall i \in \{1, \ldots, N\}$, $\forall (j_1, \ldots, j_N) \in A$,

$$P^i_{(j_1,\ldots,j_N)}\left(h^i_1, \ldots, h^i_v\right) \leq P^i_{(j_1,\ldots,j_N)}\left(f_1, \ldots, f_v\right), \qquad \forall \left(f_1, \ldots, f_v\right) \in U_f.$$

Therefore, $\forall i \in \{1, \ldots, N\}$, $\forall \left(\mathbf{x}^{-i}, \mathbf{u}^i\right) \in S$,

$$\min_{\tilde{\mathbf{P}} \in U} \pi_i\left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right) = \pi_i\left(\mathbf{Q}; \mathbf{x}^{-i}, \mathbf{u}^i\right).$$

By Relation (4.5), $\left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right) \in S$ is an equilibrium of the robust finite game iff, $\forall i \in \{1, \ldots, N\}$,

$$\mathbf{x}^i \in \arg\max_{\mathbf{u}^i \in S_{a_i}} \left[\min_{\tilde{\mathbf{P}} \in U} \pi_i\left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i\right)\right] = \arg\max_{\mathbf{u}^i \in S_{a_i}} \left[\pi_i\left(\mathbf{Q}; \mathbf{x}^{-i}, \mathbf{u}^i\right)\right].$$

□

Let us give an example of an application of Theorem 4.4.2. For $i \in \{1, \ldots, N\}$, let $I(i)^+$ and $I(i)^-$ form a partition of $\{1, \ldots, v\}$. Consider U given as in Theorem 4.4.2, with the function $\mathbf{P}$ defined as follows. $\forall i \in \{1, \ldots, N\}$ and $\forall (j_1, \ldots, j_N) \in A$,

$$\begin{aligned} P^i_{(j_1,\ldots,j_N)}\left(f_1, \ldots, f_v\right) &= \sum_{\ell \in I(i)^+} \gamma^{i,\ell}_{(j_1,\ldots,j_N)} f_\ell - \sum_{\ell \in I(i)^-} \gamma^{i,\ell}_{(j_1,\ldots,j_N)} f_\ell \\ \gamma^{i,\ell}_{(j_1,\ldots,j_N)} &\geq 0, \qquad \ell = 1, \ldots, v. \end{aligned}$$

Then

$$Q^i_{(j_1,\ldots,j_N)} = \sum_{\ell \in I(i)^+} \gamma^{i,\ell}_{(j_1,\ldots,j_N)} \underline{f_\ell} - \sum_{\ell \in I(i)^-} \gamma^{i,\ell}_{(j_1,\ldots,j_N)} \overline{f_\ell}.$$

4.5 Comparison of Robust and Bayesian Finite Games

Having established, in Section 4.4, a computation method for identifying equilibria of robust finite games without private information, in this section, using illustrative examples, we compare properties of these robust games with those of their nominal-game counterparts. By Equation (4.2), each nominal game we present is in fact equivalent to the Bayesian game that assigns a symmetric distribution to the uncertainty set in the corresponding robust game. Thus, our comparisons can be said to be between robust games and these corresponding Bayesian games.

In this same vein of comparison, turning our attention to a notion of symmetry unrelated to the symmetry of probability distributions, we end this section by discussing symmetric robust games, i.e., those in which the players are indistinguishable with respect to the game structure. We prove the existence of symmetric, robust-optimization equilibria in these games, thereby establishing a result analogous to Nash's existence theorem for symmetric Nash equilibria of symmetric, complete-information, finite games [117].

4.5.1 Equilibria Sets Are Generally Not Equivalent

The set of equilibria of a robust finite game and that of its nominal counterpart, e.g., the Bayesian game which assigns a symmetric distribution to the uncertainty set, may be disjoint. For example, consider the two-player inspection game presented in Example 1 in Section 4.2.5, with

$$\begin{array}{llll} \tilde{g} \in [8,12], & \tilde{v} \in [16,24], & \tilde{h} \in [4,6], & w = 15, \\ \check{g} = 10, & \check{v} \in 20, & \check{h} = 5. & \end{array}$$

The nominal version of the game has payoff matrix

$$\begin{pmatrix} (0,-\check{h}) & (w,-w) \\ (w-\check{g},\check{v}-w-\check{h}) & (w-\check{g},\check{v}-w) \end{pmatrix} = \begin{pmatrix} (0,-5) & (15,-15) \\ (5,0) & (5,5) \end{pmatrix}.$$

For the values given above, the nominal game has a unique equilibrium, in which the employee shirks with probability 1/3 and the employer inspects with probability 2/3. In contrast, by Theorem 4.4.2, the robust game is equivalent to the complete-information inspection game with payoff matrix

$$\begin{pmatrix} (0,-\overline{h}) & (w,-w) \\ (w-\overline{g},\underline{v}-w-\overline{h}) & (w-\overline{g},\underline{v}-w) \end{pmatrix} = \begin{pmatrix} (0,-6) & (15,-15) \\ (3,-5) & (3,1) \end{pmatrix}.$$

Thus, the robust game has a different, unique equilibrium, in which the employee shirks with probability 2/5 and the employer inspects with probability 4/5.

It is not surprising that the worker would shirk with higher probability and the employer would inspect with higher probability in the robust game than in the nominal game (i.e., in the Bayesian game assigning a symmetric distribution over the uncertainty set). Indeed, in moving from the average parameter values, as used in the nominal game, to the worst-case parameter values, as used in the robust game, the employee's opportunity cost of working increases, and the employer's cost of inspecting increases. As the employee's opportunity cost of working increases, the employer expects that the employee will be less willing to work. In order to make the employee indifferent between shirking and working, the employer must therefore be more prone to inspect, despite her higher inspection cost. Conversely, as the employer's cost of inspecting increases, the employee expects that the employer will be less willing to inspect. In order to make the employer indifferent between inspecting and not inspecting, the employee must therefore be more prone to shirk.

4.5.2 Sizes of Sets of Equilibria

The set of equilibria of a robust finite game may be smaller or larger than that of the corresponding Bayesian game assigning a symmetric distribution over the uncertainty set. For an extreme example in which the set of equilibria of a robust finite game is smaller than that of the nominal-game counterpart, consider the robust game without private information and with payoff uncertainty set

$$\left\{ \begin{pmatrix} (2,\tilde{f}) & (\tilde{f},2) \\ (\tilde{f},2) & (2,\tilde{f}) \end{pmatrix} \;\middle|\; \tilde{f} \in [0,4] \right\}.$$

Consider the nominal version of the game in which $\tilde{f} = \check{f} = 2$ is commonly known with certainty by the players. In this game, all pairs of mixed strategies for the two players are Nash equilibria. In contrast, by Theorem 4.4.2, the robust game is equivalent to the complete-information game with payoff matrix

$$\begin{pmatrix} (2,0) & (0,2) \\ (0,2) & (2,0) \end{pmatrix},$$

i.e., is equivalent to the classical, complete-information game of matching pennies (see, for example, [54]), and therefore has a unique equilibrium. In moving from the robust game to its Bayesian counterpart, the set of equilibria shrinks, because the payoff uncertainty results in reduced indifference, by each player, between his two actions.

Conversely, for an equally extreme example in which the set of equilibria of a robust finite game is larger than that of the corresponding nominal game, consider the robust game without private information and with payoff uncertainty set

$$\left\{ \begin{pmatrix} (\tilde{f}_1, \tilde{f}_2) & (\tilde{f}_2, \tilde{f}_1) \\ (\tilde{f}_2, \tilde{f}_1) & (\tilde{f}_1, \tilde{f}_2) \end{pmatrix} \middle| \, (\tilde{f}_1, \tilde{f}_2) \in [0,8] \times [0,4] \right\}.$$

Consider the nominal version of the game in which $(\tilde{f}_1, \tilde{f}_2) = (\check{f}_1, \check{f}_2) = (4,2)$ is commonly known with certainty by the players. This nominal game is now equivalent to the complete-information game of matching pennies and therefore has a unique equilibrium. In contrast, by Theorem 4.4.2, the robust game is equivalent to the complete-information game with payoff matrix

$$\begin{pmatrix} (0,0) & (0,0) \\ (0,0) & (0,0) \end{pmatrix}.$$

Thus, all pairs of mixed strategies for the two players are equilibria of the robust game. In moving from the robust game to its Bayesian counterpart, the set of equilibria expands, because the payoff uncertainty results in increased indifference, by each player, between his two actions.

4.5.3 Zero-sum Becomes Non-fixed-sum under Uncertainty

In general, if we subject to uncertainty the payoff matrix in a zero-sum game, the resulting robust game will not be a fixed-sum game. For example, consider the payoff uncertainty set

$$\left\{ \begin{pmatrix} (\tilde{f}_1, -\tilde{f}_1) & (\tilde{f}_2, -\tilde{f}_2) \\ (\tilde{f}_3, -\tilde{f}_3) & (\tilde{f}_4, -\tilde{f}_4) \end{pmatrix} \;\middle|\; (\tilde{f}_1, \tilde{f}_2, \tilde{f}_3, \tilde{f}_4) \in \prod_{\ell=1}^{4} [\underline{f_\ell}, \overline{f_\ell}] \right\}.$$

In the nominal version of this game, the players commonly know with certainty that

$$(\tilde{f}_1, \tilde{f}_2, \tilde{f}_3, \tilde{f}_4) = (\check{f}_1, \check{f}_2, \check{f}_3, \check{f}_4),$$

for some $(\check{f}_1, \check{f}_2, \check{f}_3, \check{f}_4) \in \prod_{\ell=1}^{4} [\underline{f_\ell}, \overline{f_\ell}]$. In contrast, by Theorem 4.4.2, the robust game is equivalent to the complete-information game, with payoff matrix

$$\begin{pmatrix} (\underline{f_1}, -\overline{f_1}) & (\underline{f_2}, -\overline{f_2}) \\ \\ (\underline{f_3}, -\overline{f_3}) & (\underline{f_4}, -\overline{f_4}) \end{pmatrix},$$

which is not fixed-sum unless $\underline{f_\ell} - \overline{f_\ell}$ is constant for $\ell \in \{1, 2, 3, 4\}$. This result is not surprising, since the two players' worst-case perspectives need not agree.

4.5.4 Symmetric Robust Games and Symmetric Equilibria

Let us turn our attention to symmetric games and their symmetric equilibria, which comprise an important topic in the game theory literature. We end this section by showing that symmetric equilibria are guaranteed to exist in symmetric, robust finite games, just as they are in symmetric, complete-information, finite games.

Stated very generally, a symmetric game is one in which the players are indistinguishable with respect to the game's structure (action and strategy spaces, payoff functions, information, etc.). More formally, we have the following definition.

Definition 4.5.1. A finite game with complete information is said to be **symmetric** if all players have the same action space, all players' payoff functions are invariant under permutations of the other players' actions, and all players' payoff functions are equivalent. That is, a complete-information game is symmetric if

$$\begin{aligned} a_i &= a, & & i = 1, \ldots, N \\ P^i_{(j_{-i}, j_i)} &= P^{i'}_{(j_{\sigma(-i)}, j_i)}, & & i, i' = 1, \ldots, N;\ \forall (j_{-i}, j_i) \in \{1, \ldots, a\}^N;\ \forall \boldsymbol{\sigma} \in \Sigma_{N-1}, \end{aligned}$$

where

$$\begin{aligned} (j_{-i}, j) &\triangleq (j_1, \ldots, j_{i-1}, j, j_{i+1}, \ldots, j_N) \\ (j_{\boldsymbol{\sigma}(-i)}, j) &\triangleq (j_{\sigma(1)}, \ldots, j_{\sigma(i-1)}, j, j_{\sigma(i+1)}, \ldots, j_{\sigma(N)}), \end{aligned}$$

and Σ_{N-1} denotes the set of permutations of $N-1$ elements.

A tuple of players' strategies will be said to be **symmetric** if all players' strategies in the tuple are identical. In particular, a **symmetric equilibrium** refers to an equilibrium in which all players play the same strategy.

Similarly, this definition extends, as follows, to robust finite games.

Definition 4.5.2. A robust finite game with uncertainty set $U \subseteq \mathbb{R}^{N \prod_{i=1}^{N} a_i}$ and no private information is said to be **symmetric** if

$$\begin{aligned} a_i &= a, & & i = 1, \ldots, N \\ \rho_i\left(\mathbf{x}^{-i}, \mathbf{x}^i\right) &= \rho_{i'}\left(\mathbf{x}^{\boldsymbol{\sigma}(-i)}, \mathbf{x}^i\right), & & i, i' = 1, \ldots, N;\ \forall \left(\mathbf{x}^{-i}, \mathbf{x}^i\right) \in S;\ \forall \boldsymbol{\sigma} \in \Sigma_{N-1}, \end{aligned}$$

where $\left(\mathbf{x}^{\boldsymbol{\sigma}(-i)}, \mathbf{x}^i\right)$ denotes $\left(\mathbf{x}^{\sigma(1)}, \ldots, \mathbf{x}^{\sigma(i-1)}, \mathbf{x}^i, \mathbf{x}^{\sigma(i+1)}, \ldots, \mathbf{x}^{\sigma(N)}\right)$.

Accordingly, for example, the robust game presented in Example 2 of Section 4.2.5 is symmetric.

In [117], Nash proved the existence of symmetric equilibria in symmetric, finite games with complete information. We state and prove the following analogous existence result for robust games.

Theorem 4.5.1 (Existence of Symmetric Equilibria in Symmetric Robust Finite Games). *Any N-person, non-cooperative, simultaneous-move, one-shot,* symmetric *robust game, in which $N < \infty$, in which each player $i \in \{1, \ldots, N\}$ has $1 < a < \infty$ possible actions, in which the uncertainty set of payoff matrices $U \subseteq \mathbb{R}^{Na^N}$ is bounded, and in which there is no private information, has a symmetric equilibrium.*

Proof. By the definition of symmetry of a robust game, there exists a function $\rho : S \to \mathbb{R}$ such that $\rho \equiv \rho_i$, $\forall i \in \{1, \ldots, N\}$. Now define $\Phi : S \to 2^S$ as

$$\Phi(\mathbf{x}) \quad = \quad \left\{ \mathbf{y} \in S_a \;\middle|\; \mathbf{y} \in \arg\max_{\mathbf{u} \in S_a} \rho\left(\mathbf{x}^{-i}, \mathbf{u}\right) \right\},$$

where $\mathbf{x}^{-i}$ denotes the $(N-1)$-tuple $(\mathbf{x}, \mathbf{x}, \ldots, \mathbf{x})$. The N-tuple $(\mathbf{x}, \mathbf{x}, \ldots, \mathbf{x}) \in S$ is a symmetric equilibrium of the robust game iff $\mathbf{x}$ is a fixed point of Φ. From an argument paralleling that given in the proof of Theorem 4.3.2, it follows that Φ satisfies Kakutani's Fixed Point Theorem. □

Symmetric games with incomplete information may be of particular interest for two reasons. First, incomplete-information games, in which the players are indistinguishable with respect to the game structure, may be especially amenable to the common prior assumption in Harsanyi's model and to its analog, the assumption of a common uncertainty set, in our robust game model. Second, the multilinear system formulations for symmetric equilibria of symmetric, robust finite games are smaller, by a factor of N, than those for the general equilibria of these games. Indeed, in systems (4.13) and (4.14), if we replace $\mathbf{x}^i$, $i \in \{1, \ldots, N\}$, by the single $a \times 1$ vector variable $\mathbf{x}$, subsequent elimination of redundancies then reduces the number of variables and constraints in these systems by a factor of N. Thus, we may be able to compute symmetric equilibria of symmetric, robust finite games more quickly and accurately, and with less computational effort, than we can compute the general equilibria of these games.

4.6 Robust Games with Private Information

In the preceding sections, we proposed a robust optimization approach and a corresponding distribution-free equilibrium concept for modeling incomplete-information games. We proved existence and computation results. Until now, we have focused on incomplete-information games without private information. In this section, we extend our discussion to the general case, involving potentially private information.

4.6.1 Extension of Model

As in the preceding sections of this chapter of the thesis, consider an N-person, incomplete-information game, in which player $i \in \{1, \ldots, N\}$ has $a_i < \infty$ possible actions, and in which each player is in some way uncertain of the multi-dimensional payoff matrix $\tilde{\mathbf{P}}$ that parameterizes the expected payoff vector function $\boldsymbol{\pi}$. Suppose that each player may have private information about $\tilde{\mathbf{P}}$ or about the other players' beliefs. For each player $i \in \{1, \ldots, N\}$, his potentially private information may be encoded in his "type" θ_i. Since the information is potentially private, player i may be uncertain of the type $\theta_{i'}$ of player i', $i' \neq i$. Let U denote, as before, a set of possible payoff matrices $\tilde{\mathbf{P}}$. Let Θ_i denote the set of possible types of player $i \in \{1, \ldots, N\}$, and $\Theta = \prod_{i=1}^{N} \Theta_i$.

In using separate notation for the unknown payoff parameters $\tilde{\mathbf{P}}$ and the players' types $\boldsymbol{\theta}$, we make explicit the difference between the actual payoff parameters and the players' beliefs about these parameters and about the other players' convictions. In addition, this notation allows us to very clearly address the situation in which players may both possess private information and yet still be uncertain of the parameters affecting their own payoffs. In fact, the model we propose in this section is sufficiently flexible to simultaneously capture the case of no private information (the Θ_i are singletons, $\forall i \in \{1, \ldots, N\}$), the differential information setting involving all-but-self uncertainty (each $\theta_i \in \Theta_i$ is consistent with only a single $\tilde{\mathbf{P}}^i$), and the aforementioned differential information case in which agents may possess private information, while also being uncertain of both their own and others' payoff functions.

In the same spirit as does Harsanyi, we assume that the players commonly know a "prior" set $V \subseteq U \times \Theta$ of realizable tuples of payoff parameters and type vectors. While Harsanyi furthermore assumes, in terms of our notation, that the players commonly know a distribution over this set V, we assume that the players lack such distributional information or have chosen not to use it. Player i's type θ_i induces the subset $V_i(\theta_i)$ of V consistent with θ_i,

$$V_i(\theta_i) \quad = \quad \{(\mathbf{P}, \boldsymbol{\theta}_{-i}, \theta_i) \in V\}\,.$$

That is, $V_i(\theta_i)$ gives the set of tuples of payoff matrices and type vectors that player i, when he is of type θ_i, believes are possible. As does Harsanyi, throughout the remainder of this section, we require that $\cap_{i=1}^N V_i(\theta_i) \neq \emptyset$, $\forall \boldsymbol{\theta} \in \Theta$ such that $\{(\mathbf{P}, \boldsymbol{\theta}) \in V\} \neq \emptyset$, and that the true payoff matrix $\tilde{\mathbf{P}}$ belongs to the projection of $\cap_{i=1}^N V_i(\theta_i)$ onto U. The first requirement ensures that the players' beliefs are consistent, and implies that $\boldsymbol{\theta}$ belongs to the projection of $\cap_{i=1}^N V_i(\theta_i)$ onto Θ, i.e., that the players believe that the true type vector is possible. The second requirement ensures that the players believe that the true payoff matrix is possible.

In the private information setting, for $i \in \{1, \ldots, N\}$, player i's pure strategies are mappings from his type θ_i to his action space $\{1, \ldots, a_i\}$. His so-called behavioral strategies (see, for example, Chapter 3 of [54] for an introduction to behavioral strategies) are mappings from his type θ_i to probability distributions over his action space $\{1, \ldots, a_i\}$. More formally, we denote a behavioral strategy for player i by $\mathbf{b}^i : \Theta_i \to S_{a_i}$. That is, under behavioral strategy $\mathbf{b}^i$, if player i is of type θ_i, then he plays action $j_i \in \{1, \ldots, a_i\}$ with probability

$\mathbf{b}^i_{j_i}(\theta_i)$. Let us define the notation

$$\begin{aligned}
B_{a_i} &\triangleq \left\{\mathbf{b}^i : \Theta_i \to S_{a_i}\right\} \\
B &\triangleq \prod_{i=1}^{N} B_{a_i} \\
B_{-i} &\triangleq \prod_{\substack{i'=1 \\ i' \neq i}}^{N} B_{a_i} \\
\mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}) &\triangleq \left(\mathbf{b}^1(\theta_1), \ldots, \mathbf{b}^{i-1}(\theta_{i-1}), \mathbf{b}^{i+1}(\theta_{i+1}), \ldots, \mathbf{b}^N(\theta_N)\right) \\
\left(\mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{b}^i(\theta_i)\right) &\triangleq \left(\mathbf{b}^1(\theta_1), \ldots, \mathbf{b}^N(\theta_N)\right).
\end{aligned}$$

Recall that, in Harsanyi's model, each player seeks to optimize his average performance, i.e., his average expected payoff, where the average is taken with respect to a probability distribution over $V_i(\theta_i)$. That is, in terms of our notation, in Harsanyi's model, the set of best responses by player $i \in \{1, \ldots, N\}$, when he is of type $\theta_i \in \Theta_i$, to $\mathbf{b}^{-i}(\cdot)$ is given by

$$\arg\max_{\mathbf{u}^i \in S_{a_i}} \left(\underset{\left(\tilde{\mathbf{P}}, \boldsymbol{\theta}\right) \in V_i(\theta_i)}{E} \left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{u}^i \right) \;\middle|\; \theta_i \right] \right),$$

where the expectation is taken with respect to the conditional probability distribution induced by θ_i over V. We use the notation $\mathbf{b}^{-i}(\cdot)$ to highlight the fact that $\mathbf{b}^{-i}$ is a function. Since the best response correspondence completely determines the criterion for equilibrium, it follows that the tuple of behavioral strategies $\left(\mathbf{b}^1(\cdot), \ldots, \mathbf{b}^N(\cdot)\right) \in B$ is a Bayesian equilibrium in Harsanyi's model iff, $\forall i \in \{1, \ldots, N\}$,

$$b^i(\theta_i) \in \arg\max_{\mathbf{u}^i \in S_{a_i}} \left(\underset{\left(\tilde{\mathbf{P}}, \boldsymbol{\theta}\right) \in V_i(\theta_i)}{E} \left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{u}^i \right) \;\middle|\; \theta_i \right] \right), \qquad \forall \theta_i \in \Theta_i.$$

In contrast to Harsanyi, we assume that each player $i \in \{1, \ldots, N\}$ lacks distributional information over V and $V_i(\theta_i)$ and therefore seeks to optimize his worst-case performance, i.e., his worst-case expected payoff, where the worst case is taken with respect to $V_i(\theta_i)$. Therefore, in a robust game involving private information, the set of best responses by player $i \in \{1, \dots, N\}$, when he is of type $\theta_i \in \Theta_i$, to $\mathbf{b}^{-i}(\cdot)$ is given by the set

$$\arg\max_{\mathbf{u}^i \in S_{a_i}} \left(\inf_{(\tilde{\mathbf{P}}, \boldsymbol{\theta}) \in V_i(\theta_i)} \left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{u}^i \right) \right] \right).$$

Accordingly, the tuple of behavioral strategies $\left(\mathbf{b}^1(\cdot), \dots, \mathbf{b}^N(\cdot)\right) \in B$ is an equilibrium of the robust game with private information, i.e., is a robust-optimization equilibrium of the corresponding incomplete-information game, iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{b}^i(\theta_i) \in \arg\max_{\mathbf{u}^i \in S_{a_i}} \left(\inf_{(\tilde{\mathbf{P}}, \boldsymbol{\theta}) \in V_i(\theta_i)} \left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{u}^i \right) \right] \right), \qquad \forall \theta_i \in \Theta_i.$$

Before turning to the issue of equilibria existence, let us revisit the relation of the *ex post* equilibria of an incomplete-information game to the corresponding robust-optimization equilibria, this time in the context involving potentially private information. In any such game, the tuple of behavioral strategies $\left(\mathbf{b}^1(\cdot), \dots, \mathbf{b}^N(\cdot)\right) \in B$ is an *ex post* equilibrium of the incomplete-information game, iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{b}^i(\theta_i) \in \arg\max_{\mathbf{u}^i \in S_{a_i}} \left(\left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{u}^i \right) \right] \right), \qquad \forall \theta_i \in \Theta_i;\ \forall \left(\tilde{\mathbf{P}}, \boldsymbol{\theta} \right) \in V_i(\theta_i).$$

By a proof analogous to that of Lemma 4.2.1, we may extend the result of that lemma to the general case involving potentially private information.

Lemma 4.6.1. *The set of ex post equilibria of an incomplete-information game is contained in the corresponding set of robust-optimization equilibria.*

4.6.2 Existence of Equilibria

We will now extend our existence result from Section 4.3, in which we considered robust finite games without private information, to general robust finite games. Let us start by considering such games in which all of the players' type spaces are finite, i.e., $\forall i \in \{1, \dots, N\}$, $\Theta_i = \{1, \dots, t_i\}$, where $t_i < \infty$. Recall that player i's pure strategies are mappings from Θ_i

to $\{1, \ldots, a_i\}$. Then, the set of player i's pure strategies is simply $\{1, \ldots, a_i\}^{t_i}$. Similarly, player i's behavioral strategies can be encoded as $a_i \times t_i$ matrices, where column $\ell \in \Theta_i$ gives player i's randomization over his action space when he is of type $\theta_i = \ell$. More precisely,

$$B_{a_i} = \left\{ \mathbf{X} \in \mathbb{R}^{a_i \times t_i} \mid \mathbf{X}_\ell \in S_{a_i}, \; \ell \in \Theta_i \right\},$$

where $\mathbf{X}_\ell$ denotes the ℓ^{th} column of the matrix $\mathbf{X}$. Let us define the additional shorthands

$$\begin{aligned} \mathbf{X}^{-i}_{\boldsymbol{\theta}_{-i}} &\triangleq \left(\mathbf{X}^1_{\theta_1}, \ldots, \mathbf{X}^{i-1}_{\theta_{i-1}}, \mathbf{X}^{i+1}_{\theta_{i+1}}, \ldots, \mathbf{X}^N_{\theta_N} \right) \\ \tau_i \left(\theta_i; \mathbf{X}^{-i}, \mathbf{x}^i \right) &\triangleq \inf_{\left(\tilde{\mathbf{P}}, \boldsymbol{\theta}\right) \in V_i(\theta_i)} \left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{X}^{-i}_{\boldsymbol{\theta}_{-i}}, \mathbf{x}^i \right) \right], \end{aligned}$$

where $\mathbf{X}^i_{\theta_i}$ denotes the θ_i^{th} column of the matrix $\mathbf{X}^i$. That is, τ_i denotes player i's worst-case expected payoff function.

Theorem 4.6.1. *Consider an N-person, non-cooperative, simultaneous-move, one-shot robust game, in which $N < \infty$, in which player $i \in \{1, \ldots, N\}$ has $1 < a_i < \infty$ possible actions, and in which the prior uncertainty set of payoff matrices $U \subseteq \mathbb{R}^{N \prod_{i=1}^N a_i}$ is bounded. Suppose that, $\forall i \in \{1, \ldots, N\}$, player i's type space is given by $\Theta_i = \{1, \ldots, t_i\}$, where $t_i < \infty$. Then the robust game has an equilibrium in B.*

Proof. Let us define the point-to-set mapping $\Psi : B \to 2^B$, where 2^B is the power set of B, as

$$\Psi \left(\mathbf{X}^1, \ldots, \mathbf{X}^N \right) = \left\{ \left(\mathbf{Y}^1, \ldots, \mathbf{Y}^N \right) \in B \;\middle|\; \mathbf{Y}^i_{\theta_i} \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \tau_i \left(\theta_i; \mathbf{X}^{-i}, \mathbf{u}^i \right), \; \forall i \in \{1, \ldots, N\}, \; \forall \theta_i \in \Theta_i \right\}.$$

It is obvious that $\left(\mathbf{X}^1, \ldots, \mathbf{X}^N\right)$ is a behavioral strategy equilibrium of the robust game iff it is a fixed point of Ψ. That Ψ satisfies the conditions of Kakutani's Fixed Point Theorem [79] follows from the facts that, $\forall i \in \{1, \ldots, N\}$ and $\forall \theta_i \in \Theta_i$, $\tau_i \left(\theta_i; \mathbf{X}^{-i}, \mathbf{x}^i \right)$ is continuous on $B_{-i} \times S_{a_i}$ and is concave in $\mathbf{x}^i$ over S_{a_i} for fixed $\mathbf{X}^{-i} \in B_{-i}$. The details of the proof are analogous to those in our proof of Theorem 4.3.2, and we therefore omit them. □

Having treated the case in which all of the players' type spaces are finite, let us now consider the more general case in which there may exist an $i \in \{1, \ldots, N\}$ such that $|\Theta_i| = \infty$. If player i has infinitely many types, his behavioral strategies

$$B_{a_i} = \left\{\mathbf{b}^i : \Theta_i \to S_{a_i}\right\}$$

cannot be encoded as finite matrices but are functions with infinite domains, and therefore belong to an infinite dimensional space. Kakutani's Fixed Point Theorem applies to correspondences defined over Euclidean spaces, which are, by definition, finite dimensional. Accordingly, we cannot use Kakutani's theorem to prove the existence of behavioral strategy equilibria in robust finite games in which at least one player's type space is infinite. Instead, we need a fixed point theorem that applies to Banach spaces. The following fixed point result of Bohnenblust and Karlin [23] generalizes Kakutani's theorem to Banach spaces. Before stating it, we first recall a relevant definition.

Definition 4.6.1 (as stated in Smart [148]). Let $\mathcal{S}$ and $\mathcal{T}$ be subsets of a normed space. Ψ is called a **K-mapping** of $\mathcal{S}$ into $\mathcal{T}$ if the following two conditions hold.

1. $\forall s \in \mathcal{S}$, $\Psi(s) \subseteq \mathcal{T}$, $\Psi(s) \neq \emptyset$, and $\Psi(s)$ is compact and convex.
2. The graph $\{(s,t) \mid t \in \Psi(s)\}$ is closed in $\mathcal{S} \times \mathcal{T}$.

Theorem 4.6.2 (Bohnenblust and Karlin [23], as restated in Smart [148]). *Let $\mathcal{M}$ be a closed, convex subset of a Banach space, and let Ψ be a K-mapping of $\mathcal{M}$ into a compact subset $\mathcal{M}'$ of $\mathcal{M}$. Then $\exists x \in \mathcal{M}$ such that $x \in \Psi(x)$.*

In order to apply this theorem to prove the existence of behavioral strategy equilibria in robust finite games with private information and potentially infinite type spaces,[5] we must

[5]Recall that a game is said to be finite if the number of players and the number of actions available to each player are all finite. Accordingly, it is possible for a finite game with incomplete information to involve infinite type spaces.

first establish some preliminary results. In the next two lemmas, $\forall i \in \{1, \ldots, N\}$, we consider the metric space $(B_{-i} \times S_{a_i})\,[d]$, with metric d defined as follows. $\forall\left(\mathbf{b}^{-i}(\cdot), \mathbf{x}^i\right), \left(\mathbf{f}^{-i}(\cdot), \mathbf{y}^i\right) \in B_{-i} \times S_{a_i}$,

$$
\begin{aligned}
d\left(\left(\mathbf{b}^{-i}(\cdot), \mathbf{x}^i\right), \left(\mathbf{f}^{-i}(\cdot), \mathbf{y}^i\right)\right) & \\
\triangleq \max\left\{\|\mathbf{y}^i - \mathbf{x}^i\|_\infty, \right. & \left. \max_{\substack{i' \in \{1,\ldots,N\}\setminus\{i\} \\ j_{i'} \in \{1,\ldots,a_{i'}\}}} \left[\sup_{\theta_{i'} \in \Theta_{i'}} \left| f^{i'}_{j_{i'}}(\theta_{i'}) - b^{i'}_{j_{i'}}(\theta_{i'})\right|\right]\right\}.
\end{aligned}
$$

Lemma 4.6.2. *Let $U \subseteq \mathbb{R}^{N\prod_{i=1}^N a_i}$ be bounded. Then $\forall \epsilon > 0$, $\exists \delta(\epsilon) > 0$ such that, $\forall i \in \{1, \ldots, N\}$, $\forall \theta_i \in \Theta_i$, and $\forall\left(\mathbf{b}^{-i}(\cdot), \mathbf{x}^i\right), \left(\mathbf{f}^{-i}(\cdot), \mathbf{y}^i\right) \in B_{-i} \times S_{a_i}$,*

$$
d\left(\left(\mathbf{b}^{-i}(\cdot), \mathbf{x}^i\right), \left(\mathbf{f}^{-i}(\cdot), \mathbf{y}^i\right)\right) < \delta(\epsilon)
$$

implies that, $\forall \left(\tilde{\mathbf{P}}, \boldsymbol{\theta}\right) \in V_i(\theta_i)$,

$$
\left|\pi_i\left(\tilde{\mathbf{P}}; \mathbf{f}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{y}^i\right) - \pi_i\left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{x}^i\right)\right| < \epsilon.
$$

Proof. $\forall \epsilon > 0$, consider

$$
\delta(\epsilon) = \frac{\min\{\epsilon, 1\}}{2\left(2^N - 1\right) M \prod_{i=1}^N a_i},
$$

where $1 < M < \infty$ satisfies

$$
\left|\tilde{P}^i_{(j_1,\ldots,j_N)}\right| \le M, \qquad \forall i \in \{1, \ldots, N\},\ \forall (j_1, \ldots, j_N) \in \prod_{i=1}^N \{1, \ldots, a_i\},\ \forall \tilde{\mathbf{P}} \in U.
$$

The result follows from algebraic manipulation. □

Lemma 4.6.2 immediately gives rise to the following continuity result.

Lemma 4.6.3. *Let $U \subseteq \mathbb{R}^{N\prod_{i=1}^{N} a_i}$ be bounded. Then $\forall i \in \{1, \ldots, N\}$, $\forall \theta_i \in \Theta_i$,*

$$\tau_i\left(\theta_i; \mathbf{b}^{-i}(\cdot), \mathbf{x}^i\right) \triangleq \inf_{\left(\tilde{\mathbf{P}}, \boldsymbol{\theta}\right) \in V_i(\theta_i)} \left[\pi_i\left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{x}^i\right)\right]$$

is continuous on $B_{-i} \times S_{a_i}$.

In addition, it is trivial to prove the following lemma.

Lemma 4.6.4. *$\forall i \in \{1, \ldots, N\}$, $\forall \theta_i \in \Theta_i$, and $\forall \mathbf{b}^{-i}(\cdot) \in B_{-i}$ fixed, $\tau_i\left(\theta_i; \mathbf{b}^{-i}(\cdot), \mathbf{x}^i\right)$ is concave in $\mathbf{x}^i$ over S_{a_i}.*

We may now apply Bohnenblust's and Karlin's fixed point theorem to prove the existence of behavioral strategy equilibria in robust finite games with potentially infinite type spaces.

Theorem 4.6.3 (Existence of Equilibria in Robust Finite Games). *Consider an N-person, non-cooperative, simultaneous-move, one-shot robust game, in which $N < \infty$, in which player $i \in \{1, \ldots, N\}$ has $1 < a_i < \infty$ possible actions, in which player i's type space is given by Θ_i, and in which the prior uncertainty set of payoff matrices $U \subseteq \mathbb{R}^{N\prod_{i=1}^{N} a_i}$ is bounded. This robust game has an equilibrium in B.*

Proof. We will proceed by constructing a point-to-set mapping that satisfies the conditions of Bohnenblust's and Karlin's fixed point theorem, and whose fixed points are precisely the behavioral strategy equilibria of the robust game with private information. Recall that, for a non-empty set Θ_i, the vector space of all bounded functions defined on Θ_i is a Banach space under the supremum norm (Theorem 3-2.4 of [98]). Furthermore, the direct product of finitely many Banach spaces is a Banach space (Theorem 2-4.6 of [98]). Accordingly,

$$F = \prod_{i=1}^{N} \left\{\mathbf{f}^i : \Theta_i \to \mathbb{R}^{a_i} \mid \mathbf{f}^i \text{ is bounded}\right\}$$

is a Banach space. In the notation we used to state Bohnenblust's and Karlin's fixed point theorem, take $\mathcal{M} = \mathcal{M}' = B$. B is a convex, closed, and compact subset of F.

Let us define the point-to-set mapping $\Psi : B \longrightarrow 2^B$ as

$$\Psi\left(\mathbf{b}^1(\cdot), \ldots, \mathbf{b}^N(\cdot)\right) = \left\{ \left(\mathbf{y}^1(\cdot), \ldots, \mathbf{y}^N(\cdot)\right) \in B \;\middle|\; \mathbf{y}^i(\theta_i) \in \arg\max_{\mathbf{u}^i \in S_{a_i}} \tau_i\left(\theta_i; \mathbf{b}^{-i}(\cdot), \mathbf{u}^i\right), \right.$$
$$\left. \forall i \in \{1, \ldots, N\},\ \forall \theta_i \in \Theta_i \right\}.$$

The rest of the proof follows similarly to that of Theorem 4.3.2. □

4.6.3 Computation of Equilibria

Having extended our equilibria existence result to incomplete-information games involving private information, let us now establish that one may compute these robust-optimization equilibria, when the players' type spaces are finite, via a formulation analogous to the one we gave in Section 4.4 for the case without private information.

Theorem 4.6.4. *Consider an N-person, non-cooperative, simultaneous-move, one-shot robust game, in which $N < \infty$, in which player $i \in \{1, \ldots, N\}$ has $1 < a_i < \infty$ possible actions, and in which the prior uncertainty set of payoff matrices $U \subseteq \mathbb{R}^{N \prod_{i=1}^{N} a_i}$ is bounded. Suppose $\forall i \in \{1, \ldots, N\}$, player i's type space is given by $\Theta_i = \{1, \ldots, t_i\}$, where $t_i < \infty$. Let*

$$\begin{aligned} V(\boldsymbol{\theta}) &= \{(\mathbf{P}, \boldsymbol{\theta}) \in V\} \\ T_i(\theta_i) &= \{(\boldsymbol{\theta}_{-i}, \theta_i) \in \Theta \mid V(\boldsymbol{\theta}_{-i}, \theta_i) \neq \emptyset\}, \end{aligned}$$

and $Proj(A, A')$ denote the projection of a set A onto a set A'. In addition, suppose that, $\forall i \in \{1, \ldots, N\}$, $\forall \theta_i \in \Theta_i$, $\forall \boldsymbol{\theta} \in T_i(\theta_i)$, there exists a polyhedron $U(\boldsymbol{\theta}) = Proj(V(\boldsymbol{\theta}), U)$. Then, the set of behavioral strategy equilibria of the robust game is the component-wise projection of the solution set of a system of multilinear equalities and inequalities.

Proof. Since $\forall i \in \{1, \ldots, N\}$, $t_i < \infty$,

$$B_{a_i} = \left\{\mathbf{X} \in \mathbb{R}^{a_i \times t_i} \mid \mathbf{X}_\ell \in S_{a_i},\ \ell \in \Theta_i\right\}.$$

$\left(\mathbf{X}^1, \ldots, \mathbf{X}^N\right) \in B$ is an equilibrium of this robust game iff, $\forall i \in \{1, \ldots, N\}$, $\forall \theta_i \in \Theta_i$, $\exists z^i_{\theta_i} \in \mathbb{R}$ such that $\left(\mathbf{X}^i_{\theta_i}, z^i_{\theta_i}\right)$ is a maximizer of the following robust LP, in which $\mathbf{X}^{-i}$ is regarded as data.

$$\begin{aligned}
\max_{\mathbf{X}^i_{\theta_i}, z^i_{\theta_i}} \quad & z^i_{\theta_i} \\
\text{s.t.} \quad & z^i_{\theta_i} \ \leq \ \pi_i\left(\tilde{P}; \mathbf{X}^{-i}_{\boldsymbol{\theta}_{-i}}, \mathbf{X}^i_{\theta_i}\right), \qquad && \forall \boldsymbol{\theta} \in T_i(\theta_i);\ \forall \tilde{P} \in U(\boldsymbol{\theta}) \\
& \mathbf{e}'\mathbf{X}^i_{\theta_i} \ = \ 1 \\
& \mathbf{X}^i_{\theta_i} \ \geq \ \mathbf{0},
\end{aligned}$$

where $\mathbf{e} \in \mathbb{R}^{a_i}$ is the vector of all ones. The proof follows analogously to that of Theorem 4.4.1, since $|T_i(\theta_i)| < \infty$ and $|\Theta_i| < \infty$. □

4.7 Conclusions

We make several contributions in this chapter of the thesis. We propose a novel, distribution-free model, based on robust optimization, of games with incomplete information, and we offer a corresponding distribution-free, robust-optimization equilibrium concept. We address incomplete-information games without private information as well as those involving potentially private information. Our robust optimization model of such games relaxes the assumptions of Harsanyi's Bayesian games model and simultaneously gives a notion of equilibrium that subsumes the *ex post* equilibrium concept. In addition, we prove the existence of equilibria in any such robust finite game, when the payoff uncertainty set is bounded. This existence result is in contrast to the fact that incomplete-information games need not have any *ex post* equilibria. For any robust finite game with bounded polyhedral payoff uncertainty set and finite type spaces, we formulate the set of equilibria as the dimension-reducing, component-wise projection of the solution set of a system of multilinear equations and inequalities. We suggest a computational method for approximately solving such systems and give numerical results of the implementation of this method. Furthermore, we describe a special class of robust finite games, whose equilibria are precisely those of a related complete-information game with the same number of players and the same action spaces. Using illustrative examples of robust games from this special class, we compare properties of robust finite games with those of their Bayesian-game counterparts. Moreover, we prove that symmetric equilibria exist in symmetric, robust finite games with bounded uncertainty sets.

We hope that these contributions will provide a new perspective on games with incomplete information.

Chapter 5

Robust Transportation Network Design in User- and System-Equilibrium Environments

5.1 Introduction

Network design is a central problem in the field of combinatorial optimization. Over the past fifty years, it has drawn considerable attention from the operations research, transportation planning, telecommunications engineering, computer science, and economics communities. In addition to attracting intense theoretical interest, the network design problem is of marked practical importance. Indeed, network planning and implementation, whether in the context of roadway, public utilities, or telecommunications systems, has an enormous impact on people's daily lives.

The general network design problem (NDP) is a mixed integer optimization problem, whose formulation involves a set of binary construction decision variables, a set of continuous arc capacity variables, and a set of continuous network operation variables, the last of which specifies, e.g., routing of flow on the resulting network. An NDP's objective function may depend on all three sets of variables. In the problem's most permissive instances, the constraints dictate only conservation and nonnegativity of flow. In instances with more specific requirements, the constraints may additionally prescribe that arc capacity restrictions, design expenditure limits, and equilibrium conditions on flow be satisfied.

The NDP's difficulty arises, in part, from the tradeoff between design and operational decisions; in the absence of design or routing costs, the problem may significantly simplify. For example, in the presence of linear construction expenses and the absence of arc capacity design costs and flow costs, the NDP reduces to the minimal spanning tree problem. In contrast, when there are no design expenses and the total flow costs are linear, the NDP simplifies to the shortest path problem. The steiner tree, multicommodity flow, traveling salesman, vehicle routing, and facility location problems are also special cases of the NDP.

For a thorough review of the NDP, its special forms, and solution methods, we refer the reader to Magnanti's and Wong's [102] survey paper, to the annotated bibliography of Balakrishnan, Magnanti, and Mirchandani [5], and to the references therein. In addition, the review articles of Gendron et al. [56] and Minoux [110, 111] also provide insightful overviews of the topic.

Contributors to the literature on computational approaches to the NDP have proposed numerous exact and approximate solution methods. In general, these methods belong to one of the following three categories: approaches based on Benders decomposition [9] (e.g., Geoffrion and Graves [57], Hoang [72], Magnanti et al. [100], Balakrishnan [4], and Rardin and Wolsey [134]), branch and bound algorithms (e.g., Scott [147], Boyce et al. [25], Hoang [71], LeBlanc [87], Boyce and Soberanes [26], Dionne and Florian [42], Rothengatter [139], and Los and Lardinois [95]), and Lagrangian relaxation and dual ascent methods (e.g., Rardin and Choe [133], Balakrishnan, Magnanti, and Wong [6], Hochbaum and Segev [73], and Hellstrand et al. [69]).

Although there has been progress on numerically solving the NDP, the problem is a difficult one whose general form is NP-hard. Johnson et al. [78] proved that the uncapacitated budget design problem, the version of the NDP on which we will focus in this chapter, is also NP-hard. In this version of the problem, arcs are built with limitless capacity, and the design decisions have no impact on the objective function, but their costs must satisfy a budget constraint. Wong [166] furthermore established that even approximately solving a simplified version of the budget design problem is difficult. Specifically, he focused on the problem in which all origin-destination (O-D) pairs have unit demand. In this setting, Wong proved that it is NP-hard to merely solve this specialized NDP to within a factor of $n^{1-\epsilon}$ times optimality, where n is the number of nodes in the network, and $\epsilon > 0$ is an arbitrary constant. Similarly, Plesnik [129] established that approximately solving a budget design problem with a minimax objective is also NP-hard.

Not surprisingly, given these hardness results for the NDP, other contributors to the literature have proposed heuristic methods for the problem's solution. Starting from a feasible design solution, these heuristics attempt to find local optima, or at least "good" solutions of the NDP. Many of them proceed by iteratively adding or deleting to the set of edges to be built, or by interchanging pairs of edges from the sets of arcs to be constructed and those not to be built. For example, Scott [147], Billheimer and Gray [21], Boffey and Hinxman [22], Dionne and Florian [42], Los and Lardinois [95], and Wong [166] have proposed such heuristic methods. In addition, others have offered stochastic heuristic methods based on local search, simulated annealing, evolutionary algorithms, and tabu search (e.g., Crainic et al. [36], Crainic and Gendreau [35], Randall et al. [132], and Kumar and Banerjee [84]).

5.1.1 The Network Equilibrium Problem

The objective and constraints of an NDP are defined, in part, by the criterion according to which network performance is measured. Common measures in the literature attempt to capture the total cost, possibly monetary or related to delays, experienced by all units of flow in the network. In turn, this performance depends on the behavior of the users of the system. Each such agent may act in an altruistic, cooperative manner aimed at minimizing the total cost experienced by all users of the system. Alternatively, on the opposite extreme, each agent may behave in a selfish, noncooperative fashion, routing his flow in an effort to minimize his own cost and disregarding the effect of his decisions on others.

As a result of congestion effects, the actions of one user of a network may affect the costs experienced by other users. In this way, network routing behavior may be studied in the context of game theory (see, e.g., Wardrop [164], Rosenthal [137, 138], Haurie and Marcotte [68], and Roughgarden and Tardos [143, 144]). Accordingly, the questions of how to define network routing equilibria and of whether such equilibria exist naturally arise. In a way, a flow that minimizes total system cost can be said to be at equilibrium, in that it satisfies optimality conditions. Similarly, selfish behavior may give rise to a different kind of equilibrium flow, akin to the Nash equilibrium of a noncooperative game, in which, roughly speaking, no user has incentive to unilaterally deviate by rerouting some flow.

In the context of transportation, Wardrop [164] gave two principles characterizing equilibria resulting from altruistic and selfish behavior. Dafermos and Sparrow [39] later introduced the terms "system optimal" (SO) and "user optimal" (UO) to distinguish between the two. Wardrop's user-optimal principle requires that, at user-optimal equilibrium, any path carrying a strictly positive amount of flow between a given origin and destination must be a minimum-cost path for that O-D pair. In the literature, when not explicitly specified, the term network equilibrium generally refers to a user-optimal flow solution. We use the same convention in this chapter of the thesis.

For a thorough review of the network equilibrium problem (NEP) and corresponding solution methods, we refer the interested reader to the texts by Nagurney [114] and Patriksson [123], to the survey articles by Florian and Hearn [50] and Magnanti [99], and to the references therein. In addition, since the NEP is an instance of the variational inequality (VI) problem, VI solution methods may be applied to NEPs. For a complete discussion of the history of the VI problem and its solution methods, the interested reader may consult the recent survey text by Facchinei and Pang [47], the monograph by Patriksson [124], and the references in both texts. The review article by Harker and Pang [64] and the Ph.D. thesis of Hammond [63], as well as the references therein, also provide insightful reviews of the VI problem and associated algorithms. Moreover, in Chapter 3 of this thesis, we reformulate the possibly asymmetric VI problem over a polyhedron as a single-level (and many-times continuously differentiable) optimization problem. We provide sufficient conditions for the convexity of this reformulation and prove, as a special case, that any monotone affine VI may be reformulated as a convex, linearly constrained quadratic optimization problem (LCQP). Thus, the NEP can be reformulated as a single-level optimization problem. When arc usage costs are affine and monotone with respect to arc flows, this optimization problem is a convex LCQP.

5.1.2 The Price of Anarchy

Because a selfish agent's actions may increase the cost to others more than they decrease the cost to the agent himself, such noncooperative behavior may yield a total system cost that is suboptimal. Pigou (p. 194, [128]) was perhaps the first to illustrate that selfish, noncooperative behavior may induce a strictly higher total system cost than that resulting from altruistic, cooperative actions. Koutsoupias and Papadimitriou [83] first proposed quantifying this inefficiency of UO with respect to SO. They introduced the term "price of anarchy" to refer to the inefficiency that arises from the decentralization inherent in user optimality.

Koutsoupias' and Papadimitriou's work has sparked intense activity in the computer science, operations research, transportation, and telecommunications communities. Indeed, the literature on the price of anarchy is vast and increasing, as others have since further examined the inefficiency that may arise from decentralization and competition. Roughgarden and Tardos [143, 144] and Roughgarden [141] bounded the price of anarchy of network performance in the context of uncapacitated edges and arbitrary network topology. They considered different families of separable arc cost functions, describing settings in which the flow on a given arc affects the cost on that arc alone. As a special case, Roughgarden and Tardos proved that under separable, linear cost (per unit flow) functions, the price of anarchy is no greater than $4/3$, and that this bound is tight. Chau and Sim [29] proved that this bound also applies when the arc cost functions are nonseparable but symmetric. Perakis [127] generalized the bound and the discussion to networks with nonseparable, asymmetric, and nonlinear costs, as well as to the settings of multi-period pricing and competitive supply chains. In addition, Correa, Schulz, and Stier [31] extended Roughgarden's and Tardos' results to the case of capacitated networks with separable arc cost functions.

5.1.3 Network Design under Selfish Routing and Demand Uncertainty

While the NDP is a mature area in combinatorial optimization, the analysis of instances under selfish routing or demand uncertainty, though very relevant in practice, is not fully developed. Despite the rich literatures on the network equilibrium problem and the price of anarchy, the NDP community has generally favored versions of the NDP involving centrally controlled, SO routing. It has furthermore focused on problem instances involving known, fixed demands.

For many network contexts, most notably public roadways and other transit systems, the SO flow and demand certainty assumptions are unrealistic. Indeed, in real-world transportation networks lacking a central authority directing traffic, users cannot be expected to behave in an altruistic way that benefits the community as a whole. Moreover, at least in part because of the long period of time required for a network's physical construction, during the planning phase, it is usually impossible to predict, with certainty, the network demands that will arise on the completed network.

The NDP under selfish routing is an instance of the mathematical program with equilibrium constraints (MPEC), which in turn is a generalization of the bilevel optimization problem. For an overview of MPECs, we refer the interested reader to the text by Luo, Pang, and Ralph [96]. In contrast, the NDP under SO routing has a natural formulation as a single-level optimization problem. In this way, not surprisingly, the former problem is more difficult than the latter.

This increased difficulty is due, in part, to the phenomenon illustrated by Braess' Paradox [27], which Murchland [113] introduced to the transportation community. In Braess' Paradox, the removal of an edge from a network results in an equilibrium flow solution whose total system cost is strictly less than that for the original network. Thus, Braess' seemingly counterintuitive example demonstrates that providing greater choice to selfish users, as opposed to centrally controlled users, can actually have a deleterious effect on the total system cost. Therefore, even in the network design problem with zero arc construction costs, it may be beneficial to exclude some arcs from the set of edges to be built.

The following works provide a representative, though not comprehensive, sample of the literature on the NDP under selfish routing. LeBlanc [87] considered the separable arc usage cost setting and proposed a branch and bound algorithm for the problem's solution. Others, including LeBlanc and Boyce [88] and Marcotte [104, 105], have employed solution algorithms specialized to bilevel optimization problems. In a different vein, drawing on ideas from the literature on the price of anarchy, Roughgarden [140] gave optimal inapproximability results for the NDP under selfish routing, arc construction costs of zero, and separable and nondecreasing arc usage costs. He considered the objective of minimizing the total system arc usage cost at equilibrium. Essentially, Roughgarden showed that unless $P = NP$, in this setting, there is no better approximation algorithm for this NDP than the "algorithm" of building the entire network.

Rather than addressing selfish routing considerations, others have assumed SO routing but have accounted for the network planner's uncertainty of the exact values of demands, i.e., traffic rates, on the network. Riis and Andersen [135], Lisser et al. [93], Andrade et al. [1], and Waller and Ziliaskopoulos [163], to name a few, have modeled this demand uncertainty via probability distributions and have proposed solution approaches based on stochastic optimization. In contrast, others have addressed demand uncertainty deterministically, using the robust optimization paradigm. For instance, Chekuri et al. [30], Ordóñez and Zhao [118], and Atamturk and Zhang [2] presented robust optimization models of a version of the NDP, under demand uncertainty, and in which the network planner is to determine edge capacities for a given set of edges. Chekuri et al. [30] sought the minimum-cost capacity allocation that can accommodate any realization from the uncertainty set of demands. Ordóñez and Zhao [118] treated the capacity allocation costs via a design budget and sought to minimize linear routing costs. Atamturk and Zhang [2] considered an objective arising from the sum of capacity allocation and linear routing costs. Finally, Gutiérrez et al. [61] offered a deterministic approach (though not one based on the robust optimization paradigm) to the NDP, under demand uncertainty, and in which the network planner is to determine the edge set to be built.

Although there has been some work on the NDP under selfish routing and known demands and on the NDP under SO routing and demand uncertainty, research in these areas is far from complete. For instance, among the contributions on the latter topic, robust optimization and other deterministic treatments of the problem have modeled total system routing costs as linear functions of the flow variables, and have thereby ignored the possibility of congestion effects. That is, they have modeled the cost per unit flow on an arc as a constant, rather than as a function of the flow variables. In contrast, the latter model more realistically captures conditions inherent in many real-world networks, including transportation and data systems. Furthermore, the robust optimization analyses of the NDP under demand uncertainty have focused on the capacity allocation problem, rather than the binary choice, arc construction problem. In the latter problem, the network designer must select a subset of arcs to be built, each with unlimited capacity. Finally, to our knowledge, no one has yet addressed, let alone from a robust optimization standpoint, the network design problem under *both* demand uncertainty *and* selfish routing, in the presence of congestion effects.

5.1.4 Contributions and Structure of this Chapter

In this chapter of the thesis, we make several contributions.

1. We propose a novel approach, based on robust optimization, to the binary choice, arc construction NDP, under demand uncertainty, congestion effects, and selfish routing. Our model also addresses the problem under SO routing.

2. We offer solution methods for the resulting robust NDP. In particular, we propose a branch and bound algorithm for exactly solving instances of this problem under SO routing and for heuristically solving instances of this problem under UO routing.

Moreover, we prove that the optimal solution to the robust NDP under SO routing performs, in the UO setting, within a factor of the price of anarchy times the optimal performance. In addition, we present conditions under which the robust NDP reduces to a nominal counterpart. We also characterize settings in which the robust NDP, itself a multilevel optimization problem, reduces to a single-level quadratic optimization problem.

3. Our branch and bound algorithm comprises the first constructive use of the price of anarchy, which has previously been employed only in a descriptive, rather than a prescriptive manner. Specifically, the algorithm computes upper bounds, based on the price of anarchy, in an effort to prune the branch and bound tree. In fact, these upper bounds illustrate the more general point that the use of approximation algorithms within branch and bound schemes may allow for reduced computational effort in these schemes.

4. We observe counterintuitive behavior, not yet noted in the literature, of costs at equilibrium with respect to changes in traffic demands on the network. The examples we present are analogous to Braess' Paradox [27] and illustrate that an increase in traffic demands on a network may yield a strict decrease in the costs at equilibrium.

5. Finally, we establish convexity and monotonicity properties of functions relating to the worst-case performance of a given network design decision. These properties are central to understanding the solution methods we propose and to appreciating the relative levels of difficulty among the SO and UO versions of the robust NDP and their nominal counterparts.

The structure of this chapter is as follows. In Section 5.2, we present our robust optimization model of the NDP under demand uncertainty and congestion effects, and under either SO or UO routing. We discuss, in Section 5.3, convexity and monotonicity properties of and reformulations related to these robust NDP problems. We give conditions under which the robust NDP reduces to a nominal counterpart, and we present examples of counterintuitive behavior of costs with respect to traffic demands. Motivated by the fact that the robust NDP under SO routing may be easier to solve than the same problem under UO routing, we prove, in Section 5.4, that the former problem provides a price-of-anarchy-approximate solution to the latter problem. In Section 5.5, we propose our branch and bound algorithm for the robust NDP. Finally, in Section 5.6, we present conditions under which the robust NDP under UO routing reduces to a single-level, nonconvex, quadratically constrained linear optimization problem (QCLP).

5.1.5 Notation

In addition to the notation conventions outlined in Section 1.3, throughout this chapter, $\mathbf{f}$ and $\mathbf{F}$ will denote flow vectors over the space of arcs and paths, respectively.

5.2 Formulation of the Robust Network Design Problem

In this section, we propose a distribution-free, robust optimization model for the binary choice, arc construction NDP, under demand uncertainty and congestion effects. We address both the SO and UO versions of this problem, giving extra attention to the latter setting, in which we view selfish-routing as a non-atomic congestion game.

Before we present our model, let us first establish our notation and review some relevant, basic definitions and results from the literature.

5.2.1 Review of NDP Definitions and Classical NEP Results

Consider a network $G(V, A)$, where V and A are the sets of nodes and arcs, respectively. Let W denote the set of O-D pairs, P_w denote the set of paths connecting O-D pair $w \in W$, and P denote the set of all paths, i.e., $P = \cup_{w \in W} P_w$. Let s_w and t_w denote the w^{th} origin and destination, respectively. For a set S, let $|S|$ denote its cardinality. Let $\mathbf{f} \in \mathbb{R}^{|A|}$ and

$\mathbf{F} \in \mathbb{R}^{|P|}$ denote vectors of flows on arcs and paths, respectively.

To capture the cost information on the network, let $\mathbf{c}(\mathbf{f}) : \mathbb{R}^{|A|} \to \mathbb{R}^{|A|}$ denote a vector function mapping a vector of arc flows to the vector of arc costs per unit flow. In particular, $c_a(\mathbf{f})$ denotes the cost incurred by each unit of flow on arc a. Similarly, $\mathbf{C}(\mathbf{F}) : \mathbb{R}^{|P|} \to \mathbb{R}^{|P|}$ denotes a vector function mapping a vector of path flows to the vector of path costs per unit flow. That is, $C_p(\mathbf{F})$ denotes the cost incurred by each unit of flow on path p.

A vector of path flows $\mathbf{F}$ gives rise to a corresponding vector of arc flows $\mathbf{f}$ as follows

$$f_a = \sum_{\{p \in P \mid a \in p\}} F_p, \qquad a \in A, \tag{5.1}$$

where, with a slight abuse of notation, $a \in p$ denotes that path p contains arc a. Conversely, a vector of path flows $\mathbf{F}$ is consistent with a given vector of arc flows $\mathbf{f}$ if the pair of vectors satisfies relation (5.1). Note, however, that there may be multiple, distinct path flow vectors consistent with a single, given arc flow vector. Similarly, for any pair of consistent arc and path flows $\mathbf{f}$ and $\mathbf{F}$, respectively, the vector of path costs arises from the vector of arc costs according to

$$C_p(\mathbf{F}) = \sum_{a \in p} c_a(\mathbf{f}), \qquad p \in P.$$

The way in which the amount of flow on one arc or path affects the cost on another arc or path often plays an important role in the analysis of the network equilibrium problem (NEP), and therefore, of the NDP. To that end, recall the following definitions.

Definition 5.2.1. For a set $X \subseteq \mathbb{R}^n$, a vector function $\mathbf{c} : X \to \mathbb{R}^n$ is said to be **separable** if its Jacobian matrix $J\mathbf{c}(\mathbf{x})$ is diagonal, $\forall \mathbf{x} \in X$. $\mathbf{c}$ is **symmetric** if $J\mathbf{c}(\mathbf{x})$ is a symmetric matrix, $\forall \mathbf{x} \in X$. If $\mathbf{c}$ is not symmetric, it is called **asymmetric**.

Suppose that the arc set A has not yet been built, and that a network designer is to determine, through a collection of binary decisions, a subset of A to be constructed. We focus on situations in which all constructed arcs are uncapacitated. For the remainder of this section, we state definitions and results in the context of such a network whose arc set is to be determined.

Let the network planner's decision be denoted by $\mathbf{y} \in \{0,1\}^{|A|}$. That is, $y_a = 1$ if arc $a \in A$ is to be built, and $y_a = 0$ otherwise. The network planner's design yields subnetwork $G(V, A(\mathbf{y}))$, where

$$A(\mathbf{y}) \quad = \quad \{a \in A \mid y_a = 1\}.$$

Similarly, let $P(\mathbf{y})$ denote the set of constructed paths induced by arc construction decision vector $\mathbf{y}$. Formally,

$$\mathbf{P}(\mathbf{y}) \quad = \quad \{p \in P \mid y_a = 1,\ \forall a \in p\}.$$

In this way, the set of feasible flows on network $G(V, A(\mathbf{y}))$ depends both on $\mathbf{y}$ and on the vector of O-D pair demands $\mathbf{d} \in \mathbb{R}_{+}^{|W|}$, where d_w, $w \in W$, denotes the amount of flow to be routed from $s_w \in V$ to $t_w \in V$. Let $K_A(\mathbf{y}, \mathbf{d})$ and $K_P(\mathbf{y}, \mathbf{d})$ denote the set of feasible flows on the sets of arcs and paths, respectively. Specifically, $K_A(\mathbf{y}, \mathbf{d})$ is the set of $\mathbf{f} \in \mathbb{R}^{|A|}$ such that, $\forall w \in W$, $\exists \mathbf{f}^w \in \mathbb{R}^{|A|}$ satisfying

$$\begin{aligned}
&\mathbf{f} = \sum_{w \in W} \mathbf{f}^w, && \\
&f_a^w = 0, && w \in W;\ a \in A \backslash A(\mathbf{y}), \\
&\mathbf{f}^w \geq \mathbf{0}, && w \in W, \\
\sum_{\substack{\{i \in V \text{ s.t.} \\ (v,i) \in A\}}} f_{(v,i)}^w - \sum_{\substack{\{j \in V \text{ s.t.} \\ (j,v) \in A\}}} f_{(j,v)}^w = &\begin{cases} d_w, & v = s_w, \\ -d_w, & v = t_w, \\ 0, & \text{otherwise}, \end{cases} && w \in W;\ v \in V.
\end{aligned} \tag{5.2}$$

Note that $\mathbf{f}^w$ is the vector of arc flows associated solely with the w^{th} O-D pair. Similarly,

$K_P(\mathbf{y}, \mathbf{d})$ is the set of $\mathbf{F} \in \mathbb{R}^{|P|}$ satisfying

$$\begin{aligned}
\sum_{p \in P_w} F_p &= d_w, && w \in W, \\
F_p &= 0, && p \in P \backslash P(\mathbf{y}), \\
\mathbf{F} &\geq \mathbf{0}.
\end{aligned} \tag{5.3}$$

Note that, for given $\mathbf{y}$ and $\mathbf{d}$, $K_A(\mathbf{y}, \mathbf{d})$ and $K_P(\mathbf{y}, \mathbf{d})$ define polyhedra in $\mathbb{R}^{|A|}$ and $\mathbb{R}^{|P|}$, respectively.

Throughout this chapter of the thesis, we will often want to refer to the set of design solutions yielding networks in which all of the O-D pairs are connected. Stated formally, $\mathbf{y} \in \{0,1\}^{|A|}$ is connectivity feasible if $K_A(\mathbf{y}, \mathbf{e}) \neq \emptyset$, or equivalently, if $K_P(\mathbf{y}, \mathbf{e}) \neq \emptyset$, where $\mathbf{e} \in \mathbb{R}^{|W|}$ is the vector of all ones. Note that $K_A(\mathbf{y}, \mathbf{e}) \neq \emptyset$ iff $K_A(\mathbf{y}, \mathbf{d}) \neq \emptyset$, $\forall \mathbf{d} \in \mathbb{R}_+^{|W|}$, and that $K_P(\mathbf{y}, \mathbf{e}) \neq \emptyset$ iff $K_P(\mathbf{y}, \mathbf{d}) \neq \emptyset$, $\forall \mathbf{d} \in \mathbb{R}_+^{|W|}$. The reason is that, $\forall \mathbf{d} \geq \mathbf{0}$, any elements of $K_A(\mathbf{y}, \mathbf{d})$ and $K_P(\mathbf{y}, \mathbf{d})$ are scaled versions of elements of $K_A(\mathbf{y}, \mathbf{e})$ and $K_P(\mathbf{y}, \mathbf{e})$, respectively. As a shorthand, let us define

$$Y = \left\{\mathbf{y} \in \{0,1\}^{|A|} \;\middle|\; K_A(\mathbf{y}, \mathbf{e}) \neq \emptyset\right\} = \left\{\mathbf{y} \in \{0,1\}^{|A|} \;\middle|\; K_P(\mathbf{y}, \mathbf{e}) \neq \emptyset\right\}.$$

Using this notation, let us formalize the definition for the total cost in the system experienced by all units of flow and the definition for the corresponding notion of system optimality. Recall from Section 5.1 that this total system cost is a common performance measure in the NDP literature.

Definition 5.2.2. For a given $\mathbf{y} \in Y$, a given $\mathbf{d} \in \mathbb{R}_+^{|W|}$, and a given $\mathbf{f} \in K_A(\mathbf{y}, \mathbf{d})$, the quantity $\mathbf{c}(\mathbf{f})'\mathbf{f}$ is called the **total system cost** of arc flow vector $\mathbf{f}$. Similarly, for a given $\mathbf{F} \in K_P(\mathbf{y}, \mathbf{d})$, $\mathbf{C}(\mathbf{F})'\mathbf{F}$ is called the total system cost of path flow vector $\mathbf{F}$.

Definition 5.2.3. For a given $\mathbf{y} \in Y$, and a given $\mathbf{d} \in \mathbb{R}_+^{|W|}$, $\mathbf{f} \in K_A(\mathbf{y}, \mathbf{d})$ is a **system-optimal (SO)** vector of arc flows for network $G(V, A(\mathbf{y}))$ with arc costs $\mathbf{c}$, if $\mathbf{f} \in SO_A(\mathbf{y}, \mathbf{d})$, where

$$SO_A(\mathbf{y},\mathbf{d}) \;=\; \arg\min_{\mathbf{f}} \quad \mathbf{c}(\mathbf{f})'\mathbf{f} \qquad \text{s.t.} \quad \mathbf{f} \in K_A(\mathbf{y},\mathbf{d}). \tag{5.4}$$

Similarly, $\mathbf{F} \in K_P(\mathbf{y},\mathbf{d})$ is a system-optimal vector of path flows for network $G(V, A(\mathbf{y}))$ with path costs $\mathbf{C}$ if $\mathbf{F} \in SO_P(\mathbf{y},\mathbf{d})$, where

$$SO_P(\mathbf{y},\mathbf{d}) \;=\; \arg\min_{\mathbf{F}} \quad \mathbf{C}(\mathbf{F})'\mathbf{F} \qquad \text{s.t.} \quad \mathbf{F} \in K_P(\mathbf{y},\mathbf{d}). \tag{5.5}$$

The following definition formalizes Wardrop's concept of user-optimality (UO), as mentioned in Section 5.1. Essentially, a flow vector is said to be at Wardrop equilibrium if all flow travels on minimum-cost paths.

Definition 5.2.4. For a given $\mathbf{y} \in Y$, and a given $\mathbf{d} \in \mathbb{R}_+^{|W|}$, $\mathbf{F} \in K_P(\mathbf{y},\mathbf{d})$ is a **Wardrop equilibrium** vector of path flows for network $G(V, A(\mathbf{y}))$ with path costs $\mathbf{C}$ if $\forall w \in W$, $\forall p_1, p_2 \in P_w(\mathbf{y})$,

$$F_{p_1} \;>\; 0 \quad\Longrightarrow\quad C_{p_1}(\mathbf{F}) \;\leq\; C_{p_2}(\mathbf{F}). \tag{5.6}$$

Similarly, $\mathbf{f} \in K_A(\mathbf{y},\mathbf{d})$ is a Wardrop equilibrium vector of arc flows for network $G(V, A(\mathbf{y}))$ with arc costs $\mathbf{c}$ if $\exists \mathbf{F} \in K_P(\mathbf{y},\mathbf{d})$ satisfying both (5.1) and (5.6). We denote the set of Wardrop arc- and path-flow equilibria by $UO_A(\mathbf{y},\mathbf{d})$ and $UO_P(\mathbf{y},\mathbf{d})$, respectively.

For the remainder of this chapter of the thesis, unless otherwise noted, we will use the terms equilibrium, user-optimum, and Wardrop equilibrium interchangeably. However, note that Wardrop's equilibrium concept is not the only such formalization of network equilibrium. In fact, there is another version of network equilibrium, which in the absence of continuity and monotonicity of the cost functions, differs from Wardrop's definition. We compare these two notions of equilibrium in Appendix B. In this chapter of the thesis, we consider only continuous and monotone cost functions. (The one exception to this rule is Example 5.3.3 in Section 5.3.2, in which we relax the monotonicity assumption in order to illustrate an example of paradoxical behavior of costs at equilibrium). Accordingly, we appropriately invoke Wardrop's definition of equilibrium throughout this chapter.

Let us now recall the definition of monotonicity.

Definition 5.2.5. A vector function $\mathbf{c} : X \to \mathbb{R}^n$ is said to be **monotone** on X if, $\forall \mathbf{x}^1, \mathbf{x}^2 \in X$,

$$\left[\mathbf{c}\left(\mathbf{x}^1\right) - \mathbf{c}\left(\mathbf{x}^2\right)\right]' \left(\mathbf{x}^1 - \mathbf{x}^2\right) \;\geq\; 0.$$

If this inequality is strict, $\forall \mathbf{x}^1, \mathbf{x}^2 \in X$ with $\mathbf{x}^1 \neq \mathbf{x}^2$, then $\mathbf{c}$ is called **strictly monotone**.

The following theorem establishes an equivalent and perhaps easier means of verifying monotonicity.

Theorem 5.2.1 (see, e.g., Proposition 2.3.2 of [47]). *$\mathbf{c} : X \to \mathbb{R}^n$ is monotone on X iff $J\mathbf{c}(\mathbf{x}) \succeq \mathbf{0}$, $\forall \mathbf{x} \in X$. $\mathbf{c} : X \to \mathbb{R}^n$ is strictly monotone on X iff $J\mathbf{c}(\mathbf{x}) \succ \mathbf{0}$, $\forall \mathbf{x} \in X$.*

Therefore, if $\mathbf{c}(\mathbf{x})$ is affine, i.e., if $\mathbf{c}(\mathbf{x}) = \mathbf{G}\mathbf{x} + \mathbf{h}$, then $\mathbf{c}$ is monotone iff $\mathbf{G} \succeq \mathbf{0}$ and is strictly monotone iff $\mathbf{G} \succ \mathbf{0}$. In the context of network routing costs, monotonicity of the vector of arc cost functions implies that the vector of path cost functions is also monotone.

On the topic of existence and uniqueness of Wardrop equilibria, we state a central result.

Theorem 5.2.2 (see, e.g., Florian and Hearn [50]). *Consider a network $G(V, A)$ with continuous cost functions. $\forall \mathbf{y} \in Y$, $\forall \mathbf{d} \in \mathbb{R}_+^{|W|}$, there exists a Wardrop equilibrium on $G\left(V, A(\mathbf{y})\right)$. Moreover, under strict monotonicity of the vector of cost functions, the equilibrium is unique.*

Next, recall that the Wardrop equilibrium conditions are equivalent to a variational inequality (VI) problem. The VI problem $VI(K, \mathbf{C})$ is to find an $\mathbf{x}^* \in K$ such that

$$\mathbf{C}(\mathbf{x}^*)'(\mathbf{x} - \mathbf{x}^*) \;\geq\; 0, \qquad \forall \mathbf{x} \in K.$$

Theorem 5.2.3 (Smith [149], Dafermos [38]). *Consider a network $G(V, A)$ with path cost functions $\mathbf{C}$. For a given $\mathbf{y} \in Y$, and a given $\mathbf{d} \in \mathbb{R}_+^{|W|}$, path flow vector $\mathbf{F} \in K_P(\mathbf{y}, \mathbf{d})$ is a Wardrop equilibrium iff it solves $VI(K_P(\mathbf{y}, \mathbf{d}), \mathbf{C})$. Alternatively, if $\mathbf{c}$ is the vector of arc cost functions consistent with $\mathbf{C}$, arc flow vector $\mathbf{f} \in K_A(\mathbf{y}, \mathbf{d})$ is a Wardrop equilibrium iff it solves $VI(K_A(\mathbf{y}, \mathbf{d}), \mathbf{c})$.*

Lastly, to end this section, recall that the problems of determining SO and UO flow solutions are closely related, as noted, e.g., by Beckmann et al. [8] and Dafermos and Sparrow [39]. Indeed, given $G(V, A)$, $\mathbf{y} \in Y$, $\mathbf{c} : \mathbb{R}^{|A|} \to \mathbb{R}^{|A|}$, $\mathbf{C} : \mathbb{R}^{|P|} \to \mathbb{R}^{|P|}$, and $\mathbf{d} \in \mathbb{R}_+^{|W|}$, consider problems (5.4) and (5.5) of finding corresponding SO arc and path flow solutions. From the first-order optimality conditions, and from Theorem 5.2.3, a flow solution is SO iff it is UO on the same network with a different set of costs, given by the vector of marginal costs of the original network. In terms of the arc flow variables, this vector of marginal costs is given by

$$\nabla_{\mathbf{f}} \left[\mathbf{c}(\mathbf{f})'\mathbf{f} \right] \;=\; \left[J\mathbf{c}(\mathbf{f}) \right]' \mathbf{f} + \mathbf{c}(\mathbf{f}),$$

and its a^{th} element, $a \in A$, represents the rate of change of the total system cost with respect to changes in f_a. Similarly, in terms of the path flow variables, the vector of marginal costs is given by

$$\nabla_{\mathbf{F}} \left[\mathbf{C}(\mathbf{F})'\mathbf{F} \right] \;=\; \left[J\mathbf{C}(\mathbf{F}) \right]' \mathbf{F} + \mathbf{C}(\mathbf{F}),$$

and its p^{th} element, $p \in P$, represents the rate of change of the total system cost with respect to changes in F_p. Conversely, when the vector of arc cost functions is affine and symmetric, a flow is UO iff it is twice the SO flow solution on the same network with the same costs, but with demands $\mathbf{d}/2$.

5.2.2 Modeling the NDP under Demand Uncertainty

Having established our notation, given formal definitions of system- and user-optimality, and reviewed related theoretical results from the literature, in this section, we formalize our model of the network design problem under demand uncertainty. In some parts of this formalization, we can, without loss of clarity, develop the discussion from only one of either the arc-flow or path-flow perspectives. In such contexts, we let the reader extrapolate the analogous treatment from the alternate perspective.

In order to properly frame our robust network design model, let us recall the corresponding nominal version of the problem. That is, consider the NDP in which a network planner is given a virtual network $G(V, A)$, a set W of O-D pairs, and a vector of known demands $\check{\mathbf{d}} \in \mathbb{R}_+^{|W|}$ for these O-D pairs. The planner wishes to determine which arcs to build, each with unlimited capacity. His choice is denoted by decision vector $\mathbf{y} \in \{0,1\}^{|A|}$.

We first discuss the constraints restricting the network planner's decision. Suppose that, in building arc a, he incurs a cost of b_a, and that he has a construction budget of B. That is, $\mathbf{y}$ must satisfy the budget constraint $\mathbf{b}'\mathbf{y} \leq B$. In addition, it must yield a network $G\left(V, A(\mathbf{y})\right)$ all of whose O-D pairs are connected. That is, $\mathbf{y}$ must belong to Y.

The network designer's objective is to build the subset of arcs which will minimize the total cost in the system, under either SO or UO routing. These two traffic routing paradigms give rise to two different versions of this nominal NDP, which we denote $NDP_{SO}\left(\check{\mathbf{d}}\right)$ and $NDP_{UO}\left(\check{\mathbf{d}}\right)$, respectively. Under SO routing, the planner's objective, as a function of $\mathbf{y}$, may be stated mathematically as $\zeta_{SO}\left(\mathbf{y}, \check{\mathbf{d}}\right)$, where

$$\zeta_{SO}\left(\mathbf{y}, \mathbf{d}\right) \;=\; \min_{\mathbf{f}} \left\{ \mathbf{c}(\mathbf{f})'\mathbf{f} \;\middle|\; \mathbf{f} \in K_A\left(\mathbf{y}, \mathbf{d}\right) \right\}.$$

That is, $\zeta_{SO}\left(\mathbf{y}, \mathbf{d}\right)$ denotes the total system cost of network $G(V, A(\mathbf{y}))$, with arc costs $\mathbf{c}$, under SO routing for demand vector $\mathbf{d}$.

Under non-uniqueness of user-optima, definition of the planner's objective in $NDP_{UO}\left(\check{\mathbf{d}}\right)$ requires some care. In general, different equilibria may give rise to different arc or path costs and therefore to different values of the total cost in the system. Since the network planner does not control the traffic flow, he therefore has no way of specifying, or perhaps even of predicting, which equilibrium will arise. Accordingly, we consider the version of the nominal problem $NDP_{UO}\left(\check{\mathbf{d}}\right)$ whose objective function is the highest total system cost induced by any equilibrium, i.e., whose objective function is $\zeta_{UO}\left(\mathbf{y},\check{\mathbf{d}}\right)$, where

$$\zeta_{UO}\left(\mathbf{y},\mathbf{d}\right) \;=\; \max_{\mathbf{f}}\left\{\mathbf{c}(\mathbf{f})'\mathbf{f} \;\middle|\; \mathbf{f}\in UO_A\left(\mathbf{y},\mathbf{d}\right)\right\}.$$

Stated mathematically, the nominal NDP problems, with construction budget B, are given as follows.

$$\begin{aligned} NDP_{SO}\left(\check{\mathbf{d}}\right) \;&:\; \min_{\mathbf{y}}\left\{\zeta_{SO}\left(\mathbf{y},\check{\mathbf{d}}\right) \;\middle|\; \mathbf{y}\in\{0,1\}^{|A|},\ \mathbf{b}'\mathbf{y}\le B,\ K_A(\mathbf{y},\mathbf{e})\neq\emptyset\right\},\\ NDP_{UO}\left(\check{\mathbf{d}}\right) \;&:\; \min_{\mathbf{y}}\left\{\zeta_{UO}\left(\mathbf{y},\check{\mathbf{d}}\right) \;\middle|\; \mathbf{y}\in\{0,1\}^{|A|},\ \mathbf{b}'\mathbf{y}\le B,\ K_A(\mathbf{y},\mathbf{e})\neq\emptyset\right\}.\end{aligned}$$

In contrast, in real-world settings, network designers are often uncertain of the exact value of the vector of demands $\tilde{\mathbf{d}}$ on the network. In our robust model of the NDP under demand uncertainty, we assume that the planner knows only an uncertainty set $D\subseteq\mathbb{R}_+^{|W|}$ of possible values that this vector $\tilde{\mathbf{d}}$ of demands may realize. We further assume that he either lacks distributional information on D or that he opts not to use potentially inaccurate distributional information. For the remainder of this chapter of the thesis, we assume that D is closed and bounded, and that the arc and path cost functions are continuous.

Recall that the robust optimization paradigm seeks to optimize worst-case performance with respect to the uncertainty (see, e.g., [11, 12, 13], [16, 17, 18], [45, 46], [150]). In the nominal versions of the problem, under the known vector of traffic demands $\check{\mathbf{d}}$, the performance of decision vector $\mathbf{y}\in\mathbb{R}^{|A|}$ is given by the total system cost under the SO flow pattern or least favorable UO flow pattern induced by $\mathbf{y}$. In this way, it is natural to define the robust counterparts of these problems by transforming the nominal objectives to measure worst-case performance with respect to the uncertain demand vector $\tilde{\mathbf{d}} \in D$. Whereas the nominal versions of the problem are parameterized by $\check{\mathbf{d}}$, the robust versions of the problem are parameterized by D. In the SO and UO versions of the problem, the objective functions of the network designer, who takes a robust approach to demand uncertainty, are thus given by

$$\begin{aligned} \tau_{SO}(D;\mathbf{y}) &= \max_{\tilde{\mathbf{d}}} \left\{ \zeta_{SO}\left(\mathbf{y},\tilde{\mathbf{d}}\right) \;\middle|\; \tilde{\mathbf{d}} \in D \right\}, \\ \tau_{UO}(D;\mathbf{y}) &= \max_{\tilde{\mathbf{d}}} \left\{ \zeta_{UO}\left(\mathbf{y},\tilde{\mathbf{d}}\right) \;\middle|\; \tilde{\mathbf{d}} \in D \right\}, \end{aligned}$$

respectively.

Denoting the robust NDPs under SO and UO routing by $RNDP_{SO}(D)$ and $RNDP_{UO}(D)$, respectively, we formally define these problems as follows.

$$\begin{aligned} RNDP_{SO}(D) &: \min_{\mathbf{y}} \left\{ \tau_{SO}(D;\mathbf{y}) \;\middle|\; \mathbf{y} \in \{0,1\}^{|A|},\ \mathbf{b}'\mathbf{y} \le B,\ K_A(\mathbf{y},\mathbf{e}) \neq \emptyset \right\}, \\ RNDP_{UO}(D) &: \min_{\mathbf{y}} \left\{ \tau_{UO}(D;\mathbf{y}) \;\middle|\; \mathbf{y} \in \{0,1\}^{|A|},\ \mathbf{b}'\mathbf{y} \le B,\ K_A(\mathbf{y},\mathbf{e}) \neq \emptyset \right\}. \end{aligned}$$

One may think of the formulation of $RNDP_{SO}(D)$ in the following way. A network planner must select a subset of arcs to construct. Only after he builds the network, an adversary selects from D a vector of the true values of the demands. The adversary's selection is maximally hostile with respect to the planner's chosen design. With the knowledge of the true values of the demands, the network planner may then system-optimally route the flow arising from this demand realization. In this way, $RNDP_{SO}(D)$ is an adjustable robust optimization problem, as defined by Ben-Tal et al. [10]. That is, some decisions in $RNDP_{SO}(D)$ must be made before, while other decisions may be made after the revelation of the true values of the uncertain parameters, i.e., the true demands.

One may interpret $RNDP_{UO}(D)$ in a similar fashion. After the network planner has determined the subset of arcs to build, an adversary selects from D a vector of demands, and a corresponding equilibrium flow vector for these demands. The adversary selects the values that are maximally hostile with respect to the network designer's decisions.

5.3 Properties of the Robust NDP

In this section, we discuss properties of the robust NDP, whose formulation we gave in the previous section. This discussion will allow for a better understanding of the level of difficulty of $RNDP_{UO}(D)$ with respect to $RNDP_{SO}(D)$ and of each robust NDP with respect to its nominal counterpart. Moreover, our analysis in this section will also provide deeper insight into the solution methods we propose in Sections 5.4 – 5.6.

To begin this discussion, recall the fundamental fact that, in order to solve any optimization problem, one must carry out two tasks. First, at a high level, one must navigate through the feasible region in order to identify an optimal solution. Second, at a lower level, one must be able to evaluate the objective function at given feasible solutions visited along the way. Moreover, in order to *efficiently*, i.e., in polynomial time, solve any optimization problem, it is necessary that one be able to efficiently evaluate the objective function at any given feasible solution.

In this section, we focus on properties of the robust NDP that relate specifically to the second of these two requirements, i.e., to the performance evaluation of given feasible arc design solution vectors. That is, we study the worst-case performance functions $\tau_{SO}(D;\mathbf{y})$ and $\tau_{UO}(D;\mathbf{y})$, defined in the previous section.

5.3.1 Convexity

To better understand the difficulty of evaluating $\tau_{SO}(D;\mathbf{y})$ and $\tau_{UO}(D;\mathbf{y})$, and to develop some insight into possible computation methods, in this section, let us examine the convexity, or lack thereof, of some related sets and of the functions $\zeta_{SO}(\mathbf{y},\mathbf{d})$ and $\zeta_{UO}(\mathbf{y},\mathbf{d})$ in $\mathbf{d}$.

Recall from Section 5.2.1 that, for a given $G(V,A)$, a given W, a given $\mathbf{y} \in \{0,1\}^{|A|}$, and a given $\mathbf{d} \in \mathbb{R}_+^{|W|}$, $K_A(\mathbf{y},\mathbf{d})$ and $K_P(\mathbf{y},\mathbf{d})$, the sets of feasible arc and path flow vectors, respectively, are bounded polyhedra. In addition, recall that monotonicity of the vector of arc cost functions implies that the vector of path cost functions is also monotone. Accordingly, if the vector of arc cost functions is monotone, it follows that the sets of user-optima, $UO_A(\mathbf{y},\mathbf{d})$ and $UO_P(\mathbf{y},\mathbf{d})$, are convex (see, e.g., Theorem 2.3.5 of [47]). Similarly, if $\mathbf{c}(\mathbf{f})'\mathbf{f}$ is a convex function of $\mathbf{f}$ (e.g., if $\mathbf{c}(\mathbf{f})$ is monotone and affine), then the sets of system-optima, $SO_A(\mathbf{y}, \mathbf{d})$ and $SO_P(\mathbf{y}, \mathbf{d})$ are convex.

However, for a given $\mathbf{y} \in \{0, 1\}^{|A|}$, the set of *tuples* of demand vectors and corresponding system-optima or user-optima need not be convex, even when the vector of arc cost functions is monotone and affine. Without loss of clarity, and for the sake of brevity, we frame this discussion in terms of the arc flow rather than the path flow variables. Stated formally, for a given $\mathbf{y}$, neither of the sets

$$\left\{ (\mathbf{d}, \mathbf{f}) \;\middle|\; \mathbf{d} \in D,\ \mathbf{f} \in SO_A(\mathbf{y}, \mathbf{d}) \right\}, \tag{5.7}$$

$$\left\{ (\mathbf{d}, \mathbf{f}) \;\middle|\; \mathbf{d} \in D,\ \mathbf{f} \in UO_A(\mathbf{y}, \mathbf{d}) \right\} \tag{5.8}$$

need be convex.

Example 5.3.1. As a means of illustrating, consider the example of $G(V, A(\mathbf{y}))$ and W pictured in Figure 5-1. In this network, there is a single O-D pair and two arcs connecting this pair. The first arc has cost function $10f_1$ and the second arc has cost function $f_2 + 100$, where f_i is the amount of flow on arc $i \in \{1, 2\}$. Thus, the vector of arc cost functions is separable. It is also continuous and strictly monotone, and therefore, for any given demand d to be routed from s to t, there is a corresponding unique SO flow and a corresponding unique UO flow (Theorem 5.2.2).

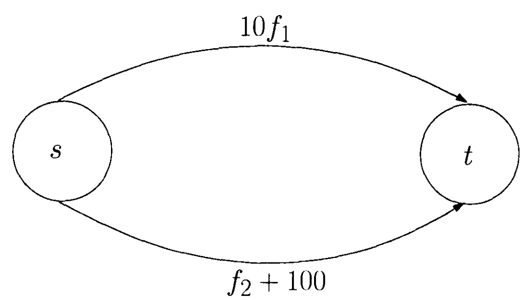

Figure 5-1: The set of tuples of demand vectors and corresponding SO or UO flows may be nonconvex.

Suppose $D = [1, \infty)$. One may easily verify that $(d, f_1, f_2) = (1, 1, 0)$ and $(d, f_1, f_2) =$

$(6, 56/11, 10/11)$ belong to set (5.7). However, their equally-weighted average does not belong to (5.7), since the unique SO flow solution for $d = 3.5$ is given by $(f_1, f_2) = (3.5, 0)$. Similarly, one may easily verify that $(d, f_1, f_2) = (1, 1, 0)$ and $(d, f_1, f_2) = (11, 111/11, 10/11)$ belong to set (5.8). However, their equally-weighted average does not belong to (5.8), since the unique UO flow solution for $d = 6$ is given by $(f_1, f_2) = (6, 0)$. Thus, sets (5.7) and (5.8) are indeed nonconvex, despite the fact that $D = [1, \infty)$ is a convex set.

In analyzing the worst-case performance functions $\tau_{SO}(D; \mathbf{y})$ and $\tau_{UO}(D; \mathbf{y})$, for a given $\mathbf{y} \in Y$, we may view these functions as maximizations of $\zeta_{SO}(\mathbf{y}, \mathbf{d})$ and $\zeta_{UO}(\mathbf{y}, \mathbf{d})$, respectively, over $\mathbf{d} \in D$. Alternatively, we may think of $\tau_{SO}(D; \mathbf{y})$ and $\tau_{UO}(D; \mathbf{y})$ as maximizations of $\mathbf{c}(\mathbf{f})'\mathbf{f}$ over $(\mathbf{d}, \mathbf{f})$ in sets (5.7) and (5.8), respectively. From either perspective, we may wish to determine whether the aforementioned objective functions, $\zeta_{SO}(\mathbf{y}, \mathbf{d})$, $\zeta_{UO}(\mathbf{y}, \mathbf{d})$, and $\mathbf{c}(\mathbf{f})'\mathbf{f}$, are convex in the corresponding variables.

With respect to the second perspective, given that sets (5.7) and (5.8) may be nonconvex even when D is convex, it is somewhat irrelevant to ask whether $\mathbf{c}(\mathbf{f})'\mathbf{f}$ is a convex function over (5.7) and (5.8). Accordingly, let us instead take the first perspective. That is, let us view $\tau_{SO}(D; \mathbf{y})$ and $\tau_{UO}(D; \mathbf{y})$, for a given $\mathbf{y} \in Y$, as maximizations of $\zeta_{SO}(\mathbf{y}, \mathbf{d})$ and $\zeta_{UO}(\mathbf{y}, \mathbf{d})$, respectively, over $\mathbf{d} \in D$. We wish to determine whether for a given $\mathbf{y}$, $\zeta_{SO}(\mathbf{y}, \mathbf{d})$ and $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ are convex functions of $\mathbf{d}$ over D.

If such convexity were to hold, it would follow that, for any given $\mathbf{y}$ and for any given closed, bounded, and convex set D, $\tau_{SO}(D; \mathbf{y})$ and $\tau_{UO}(D; \mathbf{y})$ would comprise maximizations of convex functions of $\mathbf{d}$ over the convex set D. As a result, it would follow that the values of $\tau_{SO}(D; \mathbf{y})$ and $\tau_{UO}(D; \mathbf{y})$ are realized at extreme points of D. For polyhedral D, this conclusion would, in turn, suggest a finite, though not necessarily polynomial-time, algorithm for computing $\tau_{SO}(D; \mathbf{y})$ and $\tau_{UO}(D; \mathbf{y})$ — namely, evaluating $\zeta_{SO}(\mathbf{y}, \mathbf{d})$ and $\zeta_{UO}(\mathbf{y}, \mathbf{d})$, respectively, at each of the extreme points of D. For other more sophisticated algorithms for maximizing a convex function over a convex set, see Floudas [51] and the survey paper of Floudas and Visweswaran [52].

We first consider $\zeta_{SO}(\mathbf{y}, \mathbf{d})$ and prove the following positive result relating to convexity.

Theorem 5.3.1. *Consider any network $G(V, A)$ and any set W of O-D pairs defined over V. Let D be a closed, bounded, and convex uncertainty set. Suppose that $\mathbf{c}(\mathbf{f})'\mathbf{f}$ is a convex function of $\mathbf{f}$ over $\mathbf{R}_{+}^{|A|}$. Then, for any fixed $\mathbf{y} \in Y$, $\zeta_{SO}(\mathbf{y}, \mathbf{d})$ is a convex function of $\mathbf{d}$ over D.*

Proof. Suppose D is closed, bounded, and convex, and let $\mathbf{d}^1, \mathbf{d}^2 \in D$. Let

$$\mathbf{f}^i \in \arg\min_{\mathbf{f}} \left\{ \mathbf{c}(\mathbf{f})'\mathbf{f} \;\middle|\; \mathbf{f} \in K_A\left(\mathbf{y}, \mathbf{d}^i\right) \right\}, \qquad i \in \{1, 2\}.$$

$\forall \lambda \in [0, 1]$, let

$$\mathbf{f}^3 = \arg\min_{\mathbf{f}} \left\{ \mathbf{c}(\mathbf{f})'\mathbf{f} \;\middle|\; \mathbf{f} \in K_A\left(\mathbf{y}, \lambda\mathbf{d}^1 + (1-\lambda)\mathbf{d}^2\right) \right\}. \tag{5.9}$$

Then

$$\begin{aligned} \zeta_{SO}\left(\mathbf{y}, \lambda\mathbf{d}^1 + (1-\lambda)\mathbf{d}^2\right) &= \mathbf{c}\left(\mathbf{f}^3\right)'\mathbf{f}^3 \\ &\leq \mathbf{c}\left(\lambda\mathbf{f}^1 + (1-\lambda)\mathbf{f}^2\right)'\left[\lambda\mathbf{f}^1 + (1-\lambda)\mathbf{f}^2\right] \\ &\leq \lambda\mathbf{c}\left(\mathbf{f}^1\right)'\mathbf{f}^1 + (1-\lambda)\mathbf{c}\left(\mathbf{f}^2\right)'\mathbf{f}^2 \\ &= \lambda\zeta_{SO}\left(\mathbf{y}, \mathbf{d}^1\right) + (1-\lambda)\zeta_{SO}\left(\mathbf{y}, \mathbf{d}^2\right). \end{aligned}$$

The first inequality follows from the fact that

$$\lambda\mathbf{f}^1 + (1-\lambda)\mathbf{f}^2 \in K_A\left(\mathbf{y}, \lambda\mathbf{d}^1 + (1-\lambda)\mathbf{d}^2\right),$$

but need not be a corresponding SO flow solution, i.e., need not belong to the set of optima given in (5.9). The second inequality follows from the assumption that $\mathbf{c}(\mathbf{f})'\mathbf{f}$ is a convex function of $\mathbf{f}$ over $\mathbf{R}_{+}^{|A|}$. □

Remark: When $\mathbf{c}(\mathbf{f})$ is affine and given by $\mathbf{c}(\mathbf{f}) = \mathbf{G}\mathbf{f} + \mathbf{h}$, recall that $\mathbf{c}(\mathbf{f})'\mathbf{f}$ is a convex function of $\mathbf{f}$ over $\mathbf{R}^{|A|}$ iff $\mathbf{G} \succeq \mathbf{0}$.

In contrast, we have the following negative result relating to the convexity of the function

$\zeta_{UO}(\mathbf{y}, \mathbf{d})$ in $\mathbf{d}$.

Example 5.3.2. To establish that $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ need not be a convex function of $\mathbf{d}$, consider the example of $G(V, A(\mathbf{y}))$ and W pictured in Figure 5-2. In this network, there are two O-D pairs. The first destination t_1 and second source s_2 correspond to the same node in the graph. The first O-D pair is connected by arc 1, with cost function $15f_1 + 9f_3$, where f_i denotes the flow on arc $i \in \{1, 2, 3\}$. The second O-D pair is connected by arcs 2 and 3, with cost functions $3f_2 + 5f_3 + 10$ and $5f_2 + 10f_3 + 5$, respectively. Note that the vector of arc costs is therefore of the form $\mathbf{c}(\mathbf{f}) = \mathbf{G}\mathbf{f} + \mathbf{h}$, with $\mathbf{G} \succ \mathbf{0}$. Consequently, $\forall \mathbf{d} \in \mathbb{R}_+^{|W|}$, there is a corresponding unique equilibrium flow solution.

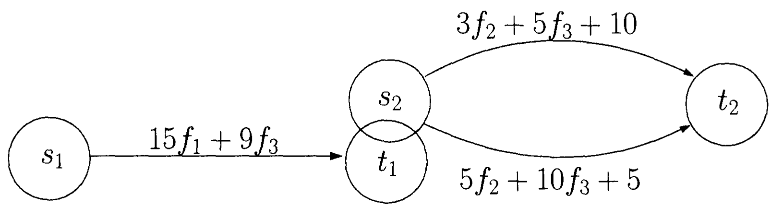

Figure 5-2: $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ may be a nonconvex function of $\mathbf{d}$ over D.

Let

$$D = \Big\{ (d_1, d_2) \;\Big|\; d_1 = 4,\; d_2 \in [0.80, 2.5] \Big\}.$$

Consider $\mathbf{d}^1 = (d_1^1, d_2^1) = (4, 0.80) \in D$. One may easily verify that $(f_1, f_2, f_3) = (4, 0, 0.80)$ is the corresponding unique equilibrium flow, yielding $\zeta_{UO}(\mathbf{y}, \mathbf{d}^1) = 279.20$. Similarly, for $\mathbf{d}^2 = (d_1^2, d_2^2) = (4, 1) \in D$, one may verify that $(f_1, f_2, f_3) = (4, 0, 1)$ is the corresponding unique equilibrium flow solution, yielding $\zeta_{UO}(\mathbf{y}, \mathbf{d}^2) = 291.00$. Finally, for $\mathbf{d}^3 = (d_1^3, d_2^3) = (4, 2.50) \in D$, the corresponding unique UO flow solution is $(f_1, f_2, f_3) = (4, 2.50, 0)$, yielding $\zeta_{UO}(\mathbf{y}, \mathbf{d}^3) = 283.75$. Since $\mathbf{d}^2$ is a convex combination of $\mathbf{d}^1$ and $\mathbf{d}^3$, it follows that $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ is a nonconvex function of $\mathbf{d}$ over D.

Thus, we have established that, under mild assumptions, for a given $\mathbf{y} \in Y$, $\zeta_{SO}(\mathbf{y}, \mathbf{d})$ is a convex function of $\mathbf{d}$ over D. In contrast, $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ need not be a convex function of $\mathbf{d}$ over D. To see that $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ furthermore need not be a concave function of $\mathbf{d}$ over D, recall Example 5.3.1. The total system costs at the unique equilibria corresponding to $d = 1$, $d = 6$, and $d = 11$ are 10, 360, and 1110, respectively. Since the total system cost at equilibrium for $d = 6$ is strictly less than the average of those for $d = 1$ and $d = 11$, it follows that $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ need not be a concave function of $\mathbf{d}$ over D.

5.3.2 Equilibrium Costs May Decrease When Demands Increase

In this section, we further discuss properties of $\tau_{UO}(D; \mathbf{y})$ and $\zeta_{UO}(\mathbf{y}, \mathbf{d})$. We again consider settings in which D is closed, bounded, and convex, and in which the vector of arc cost functions is monotone. Recall from the previous section that, under these conditions, $\tau_{UO}(D; \mathbf{y})$ need not realize its value at an extreme point of D. Our discussion in this section is motivated by the question of whether there exists a nontrivial subset of D, different from the set of extreme points, on which $\tau_{UO}(D; \mathbf{y})$ may be guaranteed to realize its value. Specifically, we examine whether $\tau_{UO}(D; \mathbf{y})$ can be guaranteed to realize its value on the boundary of D. That is, we wish to determine whether an increase in the traffic demands may yield a decrease in total system cost at equilibrium.

In the special case of a network involving a single O-D pair, we establish mild conditions under which UO total system costs are guaranteed *not* to decrease when demand increases. We prove that, under such conditions, the robust NDP therefore reduces to a nominal NDP. We give an analogous result for the SO setting. Focusing on the UO setting in the more general case, we demonstrate that, unfortunately, the same guarantee cannot be made. In particular, even under affine and monotone arc cost functions, if there is more than one O-D pair, costs may decrease as demands increase, and the worst-case vector of demands corresponding to a given $\mathbf{y}$ may in fact lie in the interior of D. This seemingly counter-intuitive behavior of costs, under UO flows and with respect to changes in traffic rates, has not yet been noted in the literature. In this way, the examples we present suggest a novel network equilibrium "paradox," analogous to Braess' Paradox [27], but stemming from changes in traffic rates rather than addition or deletion of edges from the network.

Single O-D Pair Setting: A Monotonicity Result

For the special case of a network with a single O-D pair, under mild assumptions, an increase in demand is guaranteed not to yield a decrease in total system cost. For the sake of clarity, we frame this discussion in terms of the path flow variables and path cost functions, rather than their arc-based counterparts.

To begin, let us consider the more general setting involving possibly multiple O-D pairs. We wish to examine the dependence of minimal path costs, under UO flow solutions, on the vector $\mathbf{d}$ of demands on the network. In order to guarantee existence of equilibrium flow solutions, let us assume, as we do throughout this chapter of the thesis, that the arc cost functions are continuous. Furthermore, we consider only $\mathbf{y} \in Y$. Recall from Wardrop's principle that an equilibrium flow routes all demand for O-D pair $w \in W$, on paths in $P_w(\mathbf{y})$ having minimal costs. Stated formally, at an equilibrium $\mathbf{F} \in UO_P(\mathbf{y}, \mathbf{d})$, $\forall w \in W$, $\forall p \in P_w(\mathbf{y})$,

$$C_p(\mathbf{F}) \quad \begin{cases} = \lambda_w(\mathbf{y}, \mathbf{d}), & F_p > 0, \\ \geq \lambda_w(\mathbf{y}, \mathbf{d}), & F_p = 0, \end{cases}$$

where $\lambda_w(\mathbf{y}, \mathbf{d})$ denotes the cost on used paths $p \in P_w(\mathbf{y})$, or, alternatively, the minimal path cost over $p \in P_w(\mathbf{y})$.

In the same multicommodity network setting as we consider in this chapter, Hall [62] proved the following monotonicity result with respect to the dependence of λ_w on $\mathbf{d}$. When the vector of arc cost functions is separable, and each such arc cost is positive valued, continuous, and increasing in the amount of flow on that arc, and when $\mathbf{y}$ and $d_{w'}$, $w' \neq w$, are held fixed, $\lambda_w(\mathbf{y}, \mathbf{d})$ is a nondecreasing function of d_w.[1] Hall's proof uses ideas from sensitivity analysis of convex optimization problems. For the single commodity network setting, i.e., $|W| = 1$, again involving separable, positive-valued, continuous, and nondecreasing arc cost

[1]Hall did not consider networks with variable arc sets $A(\mathbf{y})$. However, his result naturally extends to this setting. For the sake of accurately discussing his work in the context of this chapter of the thesis, we introduce this dependence of λ_w on $\mathbf{y}$.

functions, Lin, Roughgarden, and Tardos [91] gave an alternate proof that is combinatorial in nature.

We now extend Hall's result to the case of multicommodity networks with nonseparable and possibly even asymmetric arc costs. In contrast to both the proof techniques of Hall and Lin, Roughgarden, and Tardos, we base our proof on results from the theory of variational inequalities.

Before stating and proving this extension, we must clarify the setting we will consider. Under non-uniqueness of equilibria, different equilibria may give rise to different values of the vector of path cost functions and may thereby yield different values of the minimal path cost for each O-D pair. In such settings, $\lambda_w(\mathbf{y}, \mathbf{d})$ is not a well-defined function. Therefore, let us consider only settings in which the set of UO flow solutions satisfies the following property.

Definition 5.3.1 (see, e.g., Section 2.3.1 of [47]). $UO_P(\mathbf{y}, \mathbf{d})$ is **C-unique** if there exists some constant vector $\mathbf{H}$, such that, $\forall \mathbf{F} \in UO_P(\mathbf{y}, \mathbf{d})$, $\mathbf{C}(\mathbf{F}) = \mathbf{H}$.

Under equilibrium uniqueness, $UO_P(\mathbf{y}, \mathbf{d})$ is a singleton and is therefore trivially $\mathbf{C}$-unique. More interestingly, it is well-known that $UO_P(\mathbf{y}, \mathbf{d})$ is $\mathbf{C}$-unique if $\mathbf{C}$ is "pseudo monotone plus." While the definition of this characteristic is beyond the scope of this chapter, we note that it is automatically satisfied if $\mathbf{C}$ is symmetric and monotone, which occurs if $\mathbf{c}$ is symmetric and monotone. However, $\mathbf{C}$-uniqueness may also hold under less restrictive conditions that allow for asymmetry. For a more detailed discussion, we refer the interested reader to Section 2.3.1 of [47].

Under $\mathbf{C}$-uniqueness of $UO_P(\mathbf{y}, \mathbf{d})$, the function $\boldsymbol{\lambda} : Y \times \mathbb{R}_+^{|W|} \to \mathbb{R}^{|W|}$ is well-defined. Specifically, it may be expressed in terms of any $\mathbf{F} \in UO_P(\mathbf{y}, \mathbf{d})$ as follows.

$$\lambda_w(\mathbf{y}, \mathbf{d}) = \min_{p \in P_w(\mathbf{y})} C_p(\mathbf{F}). \tag{5.10}$$

We now state and prove our extension of Hall's monotonicity result.

Proposition 5.3.1. *Consider a network $G(V, A)$, with O-D pair set W and a vector $\mathbf{C}$ of* *continuous path cost functions. Suppose that* $\mathbf{C}$ *is monotone over* $\mathbb{R}_+^{|P|}$*. Furthermore, suppose that* $UO_P(\mathbf{y}, \mathbf{d})$ *is* $\mathbf{C}$*-unique.* $\forall \mathbf{y} \in Y$, $\forall w \in W$, $\forall \mathbf{d}^1 \in \mathbb{R}_+^{|W|}$, $\forall \Delta \in \mathbb{R}_+$,

$$d_q^2 = \begin{cases} d_q^1, & q \neq w, \\ d_q^1 + \Delta, & q = w, \end{cases}$$

implies $\lambda_w(\mathbf{y}, \mathbf{d}^2) \geq \lambda_w(\mathbf{y}, \mathbf{d}^1)$.

Proof. Let

$$\mathbf{F}^i \in UO_P(\mathbf{y}, \mathbf{d}^i), \qquad i \in \{1, 2\}.$$

By $\mathbf{C}$-uniqueness and definition (5.10) of $\boldsymbol{\lambda}$, it follows that

$$\mathbf{C}(\mathbf{F}^i)' \mathbf{F}^i = \boldsymbol{\lambda}(\mathbf{y}, \mathbf{d}^i)' \mathbf{d}^i, \qquad i \in \{1, 2\}.$$

However, $\forall i \neq j \in \{1, 2\}$, $\mathbf{F}^j$ need not route flow on paths with minimal costs, as induced by $C(\mathbf{F}^i)$. Accordingly,

$$\mathbf{C}(\mathbf{F}^i)' \mathbf{F}^j \geq \boldsymbol{\lambda}(\mathbf{y}, \mathbf{d}^i)' \mathbf{d}^j, \qquad (i, j) \in \{(1, 2), (2, 1)\}. \tag{5.11}$$

Since $\mathbf{C}$ is monotone over $\mathbb{R}_+^{|P|}$, by definition,

$$\begin{aligned} 0 &\leq [\mathbf{C}(\mathbf{F}^2) - \mathbf{C}(\mathbf{F}^1)]' (\mathbf{F}^2 - \mathbf{F}^1) \\ &= \mathbf{C}(\mathbf{F}^2)' \mathbf{F}^2 + \mathbf{C}(\mathbf{F}^1)' \mathbf{F}^1 - \mathbf{C}(\mathbf{F}^1)' \mathbf{F}^2 - \mathbf{C}(\mathbf{F}^2)' \mathbf{F}^1 \\ &\leq \boldsymbol{\lambda}(\mathbf{y}, \mathbf{d}^2)' \mathbf{d}^2 + \boldsymbol{\lambda}(\mathbf{y}, \mathbf{d}^1)' \mathbf{d}^1 - \boldsymbol{\lambda}(\mathbf{y}, \mathbf{d}^1)' \mathbf{d}^2 - \boldsymbol{\lambda}(\mathbf{y}, \mathbf{d}^2)' \mathbf{d}^1 \\ &= [\boldsymbol{\lambda}(\mathbf{y}, \mathbf{d}^2) - \boldsymbol{\lambda}(\mathbf{y}, \mathbf{d}^1)]' (\mathbf{d}^2 - \mathbf{d}^1) \\ &= \Delta \cdot [\lambda_w(\mathbf{y}, \mathbf{d}^2) - \lambda_w(\mathbf{y}, \mathbf{d}^1)]. \end{aligned}$$

Since $\Delta \geq 0$, this concludes the proof. □

To summarize, under continuous and monotone path costs, and under **C**-uniqueness of the corresponding set of path-flow equilibria, $\lambda_w(\mathbf{y}, \mathbf{d})$ is a nondecreasing function of d_w, when all other arguments are held fixed. In turn, this result implies the following theorem, in which D need not be connected, i.e., need not be convex.

Theorem 5.3.2. *Consider a network $G(V, A)$, with a single O-D pair and a vector $\mathbf{C}$ of continuous path cost functions. Suppose that $\mathbf{C}$ is monotone over $\mathbb{R}_+^{|P|}$. Furthermore, suppose that $UO_P(\mathbf{y}, \mathbf{d})$ is $\mathbf{C}$-unique. Then, for a given $\mathbf{y}$, $\zeta_{UO}(\mathbf{y}, d)$ is a nondecreasing function of $d \in \mathbb{R}_+$. In addition, if D is closed and bounded, then the corresponding problem $RNDP_{UO}(D)$ is equivalent to $NDP_{UO}(\bar{d})$, where*

$$\bar{d} \;=\; \max_{d \in D} d.$$

Proof. By Proposition 5.3.1, $\forall \mathbf{y} \in Y$, $\tau_{UO}(D; \mathbf{y}) = \zeta_{UO}(\mathbf{y}, \bar{d})$. □

In the next section, we show that, in more general settings, involving either multiple O-D pairs or non-monotone costs, an increase in demand may yield a strict decrease in total system cost at equilibrium. However, we first prove a result, analogous to Theorem 5.3.2, but relevant to the the case of SO routing.

Theorem 5.3.3. *Consider a network $G(V, A)$, with a single O-D pair and a vector $\mathbf{C}$ of continuous path cost functions. Suppose that $\mathbf{C}$ is monotone over $\mathbb{R}_+^{|P|}$. Furthermore, suppose that $\mathbf{C}(\mathbf{F})'\mathbf{F} \geq 0$, $\forall \mathbf{F} \in \mathbb{R}_+^{|P|}$. Then, $\forall \mathbf{y} \in Y$, $\zeta_{SO}(\mathbf{y}, d)$ is a nondecreasing function of $d \in \mathbb{R}_+$. Furthermore, if D is closed and bounded, then the corresponding problem $RNDP_{SO}(D)$ is equivalent to $NDP_{SO}(\bar{d})$, where*

$$\bar{d} \;=\; \max_{d \in D} d.$$

Proof. Consider $0 \leq d^1 \leq d^2$. Let

$$\mathbf{F}^i \;\in\; SO_P\left(\mathbf{y}, d^i\right), \qquad i \in \{1,2\}.$$

Note that, since there is only one O-D pair, $\frac{d^1}{d^2}\mathbf{F}^2 \in K_P\left(\mathbf{y}, d^1\right)$, but need not belong to $SO_P\left(\mathbf{y}, d^1\right)$. Therefore,

$$\mathbf{C}\left(\frac{d^1}{d^2}\mathbf{F}^2\right)'\left[\frac{d^1}{d^2}\mathbf{F}^2\right] \;\geq\; \mathbf{C}\left(\mathbf{F}^1\right)'\mathbf{F}^1. \tag{5.12}$$

By the monotonicity of $\mathbf{C}$ over $\mathbb{R}_+^{|P|}$,

$$\begin{aligned} 0 \;&\leq\; \left[\mathbf{C}\left(\mathbf{F}^2\right) - \mathbf{C}\left(\frac{d^1}{d^2}\mathbf{F}^2\right)\right]'\left(\mathbf{F}^2 - \frac{d^1}{d^2}\mathbf{F}^2\right) \\ &=\; \frac{d^2 - d^1}{d^2}\left[\mathbf{C}\left(\mathbf{F}^2\right) - \mathbf{C}\left(\frac{d^1}{d^2}\mathbf{F}^2\right)\right]'\mathbf{F}^2. \end{aligned}$$

Dividing through by $\frac{d^2-d^1}{d^2}$, and recalling that $d^1 \leq d^2$ and that $\mathbf{C}(\mathbf{F})'\mathbf{F} \geq 0$, $\forall \mathbf{F} \in \mathbb{R}_+^{|P|}$, we obtain

$$\frac{d^1}{d^2}\cdot\mathbf{C}\left(\frac{d^1}{d^2}\mathbf{F}^2\right)'\mathbf{F}^2 \;\leq\; \mathbf{C}\left(\frac{d^1}{d^2}\mathbf{F}^2\right)'\mathbf{F}^2 \;\leq\; \mathbf{C}\left(\mathbf{F}^2\right)'\mathbf{F}^2. \tag{5.13}$$

Relation (5.13) precludes the possibility that

$$\zeta_{SO}\left(\mathbf{y}, d^1\right) \;=\; \mathbf{C}\left(\mathbf{F}^1\right)'\mathbf{F}^1 \;>\; \mathbf{C}\left(\mathbf{F}^2\right)'\mathbf{F}^2 \;=\; \zeta_{SO}\left(\mathbf{y}, d^2\right),$$

since this inequality would contradict (5.12). Accordingly, $\forall \mathbf{y} \in Y$, $\zeta_{SO}(\mathbf{y}, d)$ is a nondecreasing function of $d \in \mathbb{R}_+$.

Finally, it follows that, $\forall \mathbf{y} \in Y$,

$$\tau_{SO}(D; \mathbf{y}) \;=\; \zeta_{SO}(\mathbf{y}, \bar{d}),$$

and, therefore, $RNDP_{SO}(D) = NDP_{SO}(\bar{d})$. □

Counter-Intuitive Examples under UO Routing

In the previous section, for a network with a single O-D pair, we gave sufficient conditions under which the equilibrium path cost and total system cost under UO routing is a nondecreasing function of d, the demand associated with the O-D pair. In this section, we show that, when there is more than one O-D pair, or when the path cost functions are not monotone (equivalently, the arc cost functions are not monotone), an increase in demands may yield a strict decrease in cost.

To begin, note that Proposition 5.3.1 says nothing about the effect that a change in d_w may have on $\lambda_{w'}(\mathbf{y}, \cdot)$, for $w, w' \in W$ with $w' \neq w$, or on the total system cost under UO routing. In addition, Proposition 5.3.1 says nothing about the effect that a change in d_w may have on O-D pair w's minimal path cost, when the arc cost functions are not monotone, even if the gradient of each arc cost function is a nonnegative vector. In the remainder of this section, we explore the nature of these relationships between $\mathbf{d}$ and the total system and minimal path costs. We show that an increase in demands *can* actually lead to a strict decrease in total system cost and in minimal path costs under UO routing. These examples may be viewed as analogous to Braess' Paradox [27] and illustrate counter-intuitive behavior not yet observed in the literature.

To that end, reconsider Example 5.3.2. In addition to illustrating that $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ need not be a convex function of $\mathbf{d}$ over D, this example also demonstrates three facts. First, recall that the move from $\mathbf{d}^2$ to $\mathbf{d}^3$ increases the level of demand d_2 corresponding to O-D pair 2 and holds the level of demand d_1 fixed. Although $\mathbf{d}^3 \geq \mathbf{d}^2$, the total system cost strictly decreases in moving from $\mathbf{d}^2$ to $\mathbf{d}^3$. Accordingly, an increase in traffic demands may yield a strict decrease in the total system cost. Second, in light of this possibility, Example 5.3.2 furthermore demonstrates that the maximal value of $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ over D (i.e., $\tau_{UO}(D; \mathbf{y})$, the worst-case total system cost under UO routing) may be realized on the interior of D. Consequently, to compute $\tau_{UO}(D; \mathbf{y})$, it is insufficient to consider only the boundary of D, let alone, in case D is convex, to consider only the extreme points of D. Third, Example 5.3.2 offers the following insight into the dependence of $\boldsymbol{\lambda}(\mathbf{y}, \mathbf{d})$ on $\mathbf{d}$. In moving from $\mathbf{d}^2$ to $\mathbf{d}^3$ in this example, note that the minimal path costs for O-D pairs 1 and 2 change from 69 to 60 and from 15 to 17.5, respectively. In this way, the increase in d_2 causes the minimal path cost for the first O-D pair to decrease significantly. That is, an increase in traffic demands can result in a decrease in equilibrium path costs.

Furthermore, let us modify Example 5.3.2 so that we instead have

$$D = \left\{ (d_1, d_2) \;\middle|\; d_1 \in [3.90, 4.05]\,,\; d_2 \in [0.80, 2.5] \right\}.$$

Consider $\mathbf{d}^4 = (d_1^4, d_2^4) = (4.05, 2.5) \in D$. The corresponding unique UO flow solution is $(4.05, 2.5, 0)$, yielding minimal path costs of 60.75 and 17.5 for O-D pairs 1 and 2, respectively, and yielding $\zeta_{UO}(\mathbf{y}, \mathbf{d}^4) = 289.79$. Thus, even if the demands for both O-D pairs simultaneously increase, as they do in moving from $\mathbf{d}^2$ to $\mathbf{d}^4$, the minimal path cost of one O-D pair and the total system cost may strictly decrease.

Furthermore, let us consider the relationship between demands and costs when the vector of arc cost functions is not monotone.

Example 5.3.3. As a means of illustrating, consider the example of $G(V, A(\mathbf{y}))$ and W pictured in Figure 5-3. In this network, there is a single O-D pair and two arcs connecting this pair. The first arc has cost function $10f_1 + 20f_2$ and the second arc has cost function $9f_1 + f_2 + 1$, where f_i is the amount of flow on arc $i \in \{1, 2\}$. Not that the vector of arc cost functions in this example is not monotone, since the matrix

$$\begin{pmatrix} 10 & 20 \\ 9 & 1 \end{pmatrix}$$

is indefinite.

For $d \leq 1/19$, there is a unique Wardrop equilibrium flow solution in which all demand is routed on arc 1 at cost $10d$ for a total system cost of $10d^2$. For $d > 1$, the unique Wardrop equilibrium flow solution has all demand on arc 2 at cost $d + 1$, for a total system cost of

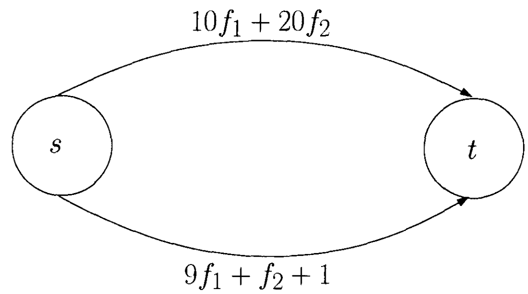

Figure 5-3: Without monotonicity, $\zeta_{UO}(\mathbf{y}, d)$ may decrease when d increases.

$d^2 + d$. For $d \in (1/19, 1]$, the following flow solutions are both Wardrop equilibria.

$$\begin{aligned} \mathbf{f} &= (f_1, f_2) = (d, 0), \\ \mathbf{f} &= (f_1, f_2) = \frac{1}{18}(19d - 1, 1 - d). \end{aligned}$$

The first of these flow solutions yields arc costs of $10d$ and a total system cost of $10d^2$. The second yields arc costs $\frac{85d+5}{9}$, for a total system cost of $\frac{85d^2+5d}{9}$.

Accordingly,

$$\zeta_{UO}(\mathbf{y}, d) = \begin{cases} 10d^2, & d \in (0, 1], \\ d^2 + d, & d \in (1, \infty). \end{cases}$$

In particular, in moving from $d = 1$ to $d = 2$, $\zeta_{UO}(\mathbf{y}, d)$ actually decreases from 10 to 6. Thus, when the vector of arc costs is not monotone, the worst-case total system cost under UO routing, as based on Wardrop's principle, can in fact decrease even when demands increase.

We digress momentarily to recall, as mentioned in Section 5.2.1, that Wardrop's principle is not the only network equilibrium concept contained in the literature. As we discuss in Appendix B, Example 5.3.3 also illustrates how Wardrop's principle may differ from other equilibrium notions when monotonicity fails to hold.

Furthermore, note that Nagurney [115] observed a phenomenon related, but not identical, to the one we have just presented. In particular, she showed that an increase in traffic demands on a network can yield a decrease in the total system "emissions" of the network. She defined the emissions to be a linear function, unrelated to the arc and path costs, of the arc flow variables. Nagurney's observation thus differs, for two major reasons, from the ones we present here. First, while the total system emissions are a linear function of the arc flow variables, the total system cost may be a nonlinear function of these variables. In particular, in our examples in this section, the total system cost is a quadratic function of the arc flows. Second, and more importantly, the coefficients defining the total system emissions are constant with respect to the flow variables and in no way influence flow costs or flow routing. In contrast, the arc cost functions that determine total system cost not only influence but essentially determine flow routing. Indeed, a selfish agent selects the path that minimizes the cost he experiences. Thus, it is perhaps more surprising that an increase in demands on the network could yield a decrease in minimal path costs and total system cost, than it is that the same increase could yield a decrease in total system emissions.

5.3.3 A Single-Level Optimization Reformulation of $\tau_{UO}(D; \mathbf{y})$

In Sections 5.3.1 and 5.3.2, we discussed several properties of $\tau_{SO}(D; \mathbf{y})$ and $\tau_{UO}(D; \mathbf{y})$. In particular, for D closed, bounded, and convex, we proved that, for a fixed $\mathbf{y} \in Y$, $\tau_{SO}(D; \mathbf{y})$ is guaranteed to realize its value at an extreme point of D. Thus, when D is furthermore polyhedral, in order to evaluate $\tau_{SO}(D; \mathbf{y})$, one could evaluate $\zeta_{SO}(\mathbf{y}, \mathbf{d})$, for $\mathbf{d}$ equal to each of the extreme points of D, and then take the maximum of these values. When D is a polyhedron and the vector of arc cost usage functions is monotone and affine, each such $\zeta_{SO}(\mathbf{y}, \mathbf{d})$ is given by a convex, linearly constrained quadratic optimization problem (LCQP).

In contrast, we showed that for D convex, closed, and bounded, $\tau_{UO}(D; \mathbf{y})$ need not realize its value at a boundary point, let alone at an extreme point of D. Thus, while we have already established at least one way of evaluating $\tau_{SO}(D; \mathbf{y})$ when D is polyhedral, we have not yet proposed a method for computing $\tau_{UO}(D; \mathbf{y})$. In this section, we offer a single-level optimization reformulation of $\tau_{UO}(D; \mathbf{y})$. From the definition of $\tau_{UO}(D; \mathbf{y})$ in Section 5.2.2, note that $\tau_{UO}(D; \mathbf{y})$ is the optimal value of a mathematical program with equilibrium constraints (MPEC). Under monotone and affine arc usage costs, for instance, this MPEC's objective is a convex function of its variables. However, the objective is to maximize, rather than minimize, this function. As such, not surprisingly, the single-level optimization reformulation we derive for $\tau_{UO}(D; \mathbf{y})$, under general costs, is nonconvex. While neither the MPEC nor the single-level formulation of $\tau_{UO}(D; \mathbf{y})$ may provide a satisfactorily practical means of exactly evaluating this function, the latter allows for greater manipulation of the problem. In particular, under monotone and affine arc usage cost functions, it yields a Lagrangian relaxation method of computing an upper bound on $\tau_{UO}(D; \mathbf{y})$.

For conciseness of notation, we frame the discussion in this section in terms of the path flow variables and path usage costs. In practice, because the number of paths in a graph may be exponential in the number of arcs in the graph, one should use the arc-based formulation, whose derivation is analogous to the path-based one we present.

The single-level optimization reformulation we propose for $\tau_{UO}(D; \mathbf{y})$ is motivated by the results of Chapter 3 of this thesis. In that chapter, we cast the variational inequality (VI) and the MPEC as special instances of the robust optimization problem. Using duality-based proof techniques from robust optimization, we gave equivalent, single-level optimization reformulations of the VI and MPEC. In particular, recall Theorem 3.2.1 from Section 3.2.1 of this thesis. From Theorems 5.2.3 and 3.2.1, it immediately follows that, $\forall \mathbf{y} \in \{0,1\}^{|A|}$, $\forall \mathbf{d} \in \mathbb{R}_+^{|W|}$, $\mathbf{F} \in UO_P(\mathbf{y}, \mathbf{d})$ iff $\exists \boldsymbol{\lambda} \in \mathbb{R}^{|W|}$ such that

$$
\begin{aligned}
&\mathbf{C}(\mathbf{F})'\mathbf{F} \leq \mathbf{d}'\boldsymbol{\lambda}, && \\
&\sum_{p \in P_w} F_p = d_w, && w \in W, \\
&F_p = 0, && p \in P \backslash P(\mathbf{y}), \\
&\mathbf{F} \geq \mathbf{0}, && \\
&\lambda_w \leq C_p(\mathbf{F}), && \forall w \in W;\ \forall p \in P_w(\mathbf{y}).
\end{aligned}
\tag{5.14}
$$

Note that the $\boldsymbol{\lambda}$ appearing in this system is precisely the vector of minimal path costs at equilibrium, as discussed in Section 5.3.2. Finally, this equivalence implies that

$$\begin{aligned}
\tau_{UO}(D;\mathbf{y}) = \max_{\tilde{\mathbf{d}},\mathbf{F},\boldsymbol{\lambda}} \quad & \mathbf{C}(\mathbf{F})'\mathbf{F} \\
\text{s.t.} \quad & (\mathbf{F},\boldsymbol{\lambda}) \text{ satisfies system (5.14), parameterized by } \tilde{\mathbf{d}} \\
& \tilde{\mathbf{d}} \in D.
\end{aligned} \tag{5.15}$$

In general, this single-level optimization reformulation (5.15) is nonconvex due, in part, to the bilinear term $\tilde{\mathbf{d}}'\boldsymbol{\lambda}$ in the first constraint of (5.14). Consequently, it may not be practical to exactly compute $\tau_{UO}(D;\mathbf{y})$ using a commercial solver.

When $\mathbf{C}(\mathbf{F})$ is affine and monotone with respect to $\mathbf{F}$, and D is a polyhedral, $\tau_{UO}(D;\mathbf{y})$ is equivalent to a nonconvex, quadratically constrained quadratic optimization problem (QCQP). The nonconvexity in this setting is due not only to the bilinear term $\tilde{\mathbf{d}}'\boldsymbol{\lambda}$, but also to the curvature of the objective function. Methods for solving indefinite QCQPs exist in the literature (see, e.g., Floudas [51]). However, corresponding software is not available to the same extent as are commercial convex optimization solvers.

Nonetheless, when $\mathbf{C}(\mathbf{F})$ is affine and monotone with respect to $\mathbf{F}$, and when D is polyhedral, one may use commercial optimization software to upper-bound QCQP (5.15) and to thereby upper-bound $\tau_{UO}(D;\mathbf{y})$. In particular, consider the following Lagrangian function $L(D,\mathbf{y};\cdot):\mathbb{R}_+ \to \mathbb{R}$ of (5.15), parameterized by D and $\mathbf{y} \in Y$ and given by

$$\begin{aligned}
L(D,\mathbf{y};\theta) = \max_{\tilde{\mathbf{d}},\mathbf{F},\boldsymbol{\lambda}} \quad & \mathbf{C}(\mathbf{F})'\mathbf{F} - \theta\left[\mathbf{C}(\mathbf{F})'\mathbf{F} - \tilde{\mathbf{d}}'\boldsymbol{\lambda}\right] \\
\text{s.t.} \quad & \sum_{p\in P_w} F_p = \tilde{d}_w, & w \in W \\
& F_p = 0, & p \in P\backslash P(\mathbf{y}) \\
& \mathbf{F} \geq \mathbf{0} \\
& \lambda_w \leq C_p(\mathbf{F}), & \forall w \in W;\ \forall p \in P_w(\mathbf{y}) \\
& \tilde{\mathbf{d}} \in D.
\end{aligned}$$

By weak duality, $\forall\theta \geq 0$, $L(D,\mathbf{y};\theta) \geq \tau_{UO}(D;\mathbf{y})$. Furthermore, this inequality may be strict, $\forall\theta \geq 0$, if there is a duality gap for this primal-dual pair. For any fixed $\theta \geq 0$, $L(D,\mathbf{y};\theta)$ is the optimal value of an LCQP with indefinite objective function. We propose to solve this LCQP via an equivalent linear, mixed integer optimization problem (MIP), whose integer variables are binary. We justify this equivalence in the following theorem.

Theorem 5.3.4. *For* $\mathbf{Q} \in \mathbb{R}^{n\times n}$, $\mathbf{r} \in \mathbb{R}^n$, $\mathbf{\Phi} \in \mathbb{R}^{m\times n}$, *and* $\boldsymbol{\psi} \in \mathbb{R}^m$, *LCQP*

$$\min_{\mathbf{x}} \left\{ \frac{1}{2}\mathbf{x}'\mathbf{Q}\mathbf{x} + \mathbf{r}'\mathbf{x} \;\middle|\; \mathbf{\Phi}\mathbf{x} \geq \boldsymbol{\psi} \right\} \tag{5.16}$$

is equivalent to the following linear, MIP with binary integer variables.

$$\begin{aligned}
\min_{\mathbf{x},\boldsymbol{\mu},\boldsymbol{\xi}} \quad & \frac{1}{2}\mathbf{r}'\mathbf{x} + \frac{1}{2}\boldsymbol{\psi}'\boldsymbol{\mu} \\
s.t. \quad & \mathbf{Q}\mathbf{x} + \mathbf{r} - \mathbf{\Phi}'\boldsymbol{\mu} \;=\; \mathbf{0} \\
& \mathbf{\Phi}\mathbf{x} \;\geq\; \boldsymbol{\psi} \\
& \mu_i \;\leq\; M\xi_i, \qquad i = 1,\ldots,m \\
& \boldsymbol{\phi}_i'\mathbf{x} - \psi_i \;\leq\; M(1-\xi_i), \qquad i = 1,\ldots,m \\
& \boldsymbol{\mu} \;\geq\; \mathbf{0},\; \boldsymbol{\xi} \;\in\; \{0,1\}^m,
\end{aligned} \tag{5.17}$$

where $\boldsymbol{\phi}_i$ *denotes the* i^{th} *row of the matrix* $\mathbf{\Phi}$, *and where* $M \in \mathbb{R}$ *is a sufficiently large, positive number.*

Proof. The Karush-Kuhn-Tucker (KKT) conditions of any linearly constrained optimization problem are necessary (see, e.g., Lemma 5.1.4 of Bazaraa and Shetty [7]). Note that $\mathbf{Q}$ may be indefinite and LCQP therefore need not be convex. Accordingly, the KKT conditions, while necessary, need not be sufficient for the optimality of a feasible solution of LCQP (5.16).

Consider any KKT point $\mathbf{x}$ of LCQP (5.16). That is, consider any $\mathbf{x} \in \mathbb{R}^n$ such that

$\exists \boldsymbol{\mu} \in \mathbb{R}^m$ satisfying

$$\mathbf{Qx} + \mathbf{r} - \boldsymbol{\Phi}'\boldsymbol{\mu} = 0 \tag{5.18}$$

$$\boldsymbol{\Phi}\mathbf{x} \geq \boldsymbol{\psi}$$

$$\boldsymbol{\mu} \geq \mathbf{0}$$

$$\mu_i \left[\boldsymbol{\phi}_i'\mathbf{x} - \psi_i\right] = 0, \qquad i = 1, \ldots, m. \tag{5.19}$$

For any such KKT point, $\mathbf{x}$, with corresponding vector $\boldsymbol{\mu}$ of multipliers, the objective value of LCQP (5.16) may be rewritten as follows. Into the objective function of LCQP (5.16), we may substitute for $\mathbf{Qx}$ from (5.18) and, subsequently, for $\mathbf{x}'\boldsymbol{\Phi}'\boldsymbol{\mu}$ from (5.19), to obtain

$$\begin{aligned} \frac{1}{2}\mathbf{x}'\mathbf{Qx} + \mathbf{r}'\mathbf{x} &= \frac{1}{2}\mathbf{x}'\left[\boldsymbol{\Phi}'\boldsymbol{\mu} - \mathbf{r}\right] + \mathbf{r}'\mathbf{x} \\ &= \frac{1}{2}\boldsymbol{\psi}'\boldsymbol{\mu} + \frac{1}{2}\mathbf{r}'\mathbf{x}, \end{aligned}$$

which is linear in $\mathbf{x}$ and $\boldsymbol{\mu}$.

Furthermore, observe that, for sufficiently large $M > 0$, the complementarity constraints (5.19) of the KKT conditions hold iff $\exists \boldsymbol{\xi} \in \mathbb{R}^m$ such that

$$\begin{aligned} \mu_i &\leq M\xi_i, & i &= 1, \ldots, m \\ \boldsymbol{\phi}_i'\mathbf{x} - \psi_i &\leq M\left(1 - \xi_i\right), & i &= 1, \ldots, m \\ \boldsymbol{\xi} &\in \{0, 1\}^m. \end{aligned}$$

If $\mathbf{x}^*$ is an optimal solution of LCQP (5.16), then it is a KKT point of LCQP (5.16). Namely, it is the KKT point yielding the highest objective value of LCQP (5.16). Therefore, $\exists \left(\boldsymbol{\mu}^*, \boldsymbol{\xi}^*\right)$ such that $\left(\mathbf{x}^*, \boldsymbol{\mu}^*, \boldsymbol{\xi}^*\right)$ is an optimal solution of MIP (5.17) with

$$\frac{1}{2}\mathbf{x}^{*\prime}\mathbf{Qx}^* + \mathbf{r}'\mathbf{x}^* = \frac{1}{2}\boldsymbol{\psi}'\boldsymbol{\mu}^* + \frac{1}{2}\mathbf{r}'\mathbf{x}^*.$$

For the reverse direction, suppose that $\left(\mathbf{x}^*, \boldsymbol{\mu}^*, \boldsymbol{\xi}^*\right)$ is an optimal solution of MIP (5.17).

Then, from the definition of MIP (5.17), $\mathbf{x}^*$ is the KKT point of LCQP (5.16) yielding the highest objective value. That is, $\mathbf{x}^*$ is an optimal solution of LCQP (5.16). □

Thus, when D is polyhedral and $\mathbf{C}(\mathbf{F})$ is affine and monotone with respect to $\mathbf{F}$, $\forall \theta \geq 0$, one may compute LCQP $L(D, \mathbf{y}; \theta)$ by solving a linear MIP with binary integer variables. This MIP thereby yields an upper bound on $\tau_{UO}(D; \mathbf{y})$. In order to derive the best such possible upper-bound on $\tau_{UO}(D; \mathbf{y})$, one would ideally like to solve $\min_{\theta \geq 0} L(D, \mathbf{y}; \theta)$. Suppose one knew the optimal solution θ^* to lie in $\left[0, \overline{\theta}\right]$. Since, $\forall \mathbf{y} \in Y$, $L(D, \mathbf{y}; \theta)$ is a convex function of θ (see, e.g., Theorem 6.3.1 of [7]), and since θ is a scalar variable, one could approximately identify θ^* using Fibonacci search over $\left[0, \overline{\theta}\right]$ (see, e.g., Section 8.1 of [7]). However, in general, one may not *a priori* know $\overline{\theta}$. In this setting, one may select an arbitrary $\overline{\theta}$ and perform Fibonacci search to obtain $\min_{\theta \in [0, \overline{\theta}]} L(D, \mathbf{y}; \theta)$, which provides an upper bound, though perhaps not the best such bound, on $\tau_{UO}(D; \mathbf{y})$.

Finally, let us briefly consider $\tau_{SO}(D; \mathbf{y})$. Recall that, for a given $\mathbf{y}$ and $\mathbf{d}$, the necessary and sufficient conditions for a flow solution to be SO comprise a VI. Accordingly, just as for $\tau_{UO}(D; \mathbf{y})$, Theorem 3.2.1 suggests a reformulation of $\tau_{SO}(D; \mathbf{y})$ as a single-level optimization problem, analogous, though not identical, to (5.15). However, this analogous reformulation also contains constraints involving a term that is bilinear in the decision variables. In this way, the single-level optimization reformulation of $\tau_{SO}(D; \mathbf{y})$ obscures this function's convexity properties, which we proved in Section 5.3.1. For this reason, we omit the explicit statement of this reformulation of $\tau_{SO}(D; \mathbf{y})$.

5.3.4 Difficulty of the Robust vs. the Nominal NDP

In Sections 5.3.1 – 5.3.3, we discussed properties and reformulations relating to the robust NDP. These properties seem to suggest that $RNDP_{UO}(D)$ is in general a more difficult problem than is $RNDP_{SO}(D)$. Let us now compare the computational demands of each of these problems with respect to those of its nominal counterpart.

In $RNDP_{SO}(D)$, recall that the worst-case performance, $\tau_{SO}(D; \mathbf{y})$ of a design decision $\mathbf{y}$ is given by the optimal value of a bilevel optimization problem. When D is closed, bounded, and convex, and when $\mathbf{c}(\mathbf{f})'\mathbf{f}$ is a convex function of $\mathbf{f}$, this bilevel optimization problem constitutes a maximization of a convex function over a convex set. In contrast, in $NDP_{SO}\left(\check{\mathbf{d}}\right)$, the performance of a design solution $\mathbf{y}$ is given by the total system cost at a SO flow for $\mathbf{y}$ and $\check{\mathbf{d}}$. Thus, evaluating the performance of $\mathbf{y}$ in the nominal setting requires only the solution of (5.4). When $\mathbf{c}(\mathbf{f})'\mathbf{f}$ is a convex function of $\mathbf{f}$, problem (5.4) is linearly constrained and convex. When $\mathbf{c}$ is furthermore affine, this problem is an LCQP.

In $RNDP_{UO}(D)$, recall that the worst-case performance, $\tau_{UO}(D;\mathbf{y})$, of a design decision vector $\mathbf{y}$ is a mathematical program with equilibrium constraints (MPEC). Even when D is polyhedral and $\mathbf{c}$ is affine and monotone with respect to the flow variables, $\tau_{UO}(D;\mathbf{y})$ represents the maximization over $\mathbf{d}$ of a function that may be neither convex nor concave in $\mathbf{d}$ and that may strictly decrease under increases in $\mathbf{d}$. Consider, in contrast, the requirements of performance evaluation in $NDP_{UO}\left(\check{\mathbf{d}}\right)$. Under the same conditions, i.e., when D is polyhedral and $\mathbf{c}$ is affine and monotone with respect to the flow variables, the set of equilibria for any given $\mathbf{y} \in Y$ is a polyhedron (see, e.g., Theorems 2.3.5 and 2.4.13 of [47]). Thus, in this setting, determining the performance of a design solution requires, at worst, the maximization of a convex function over a polyhedron.

5.4 An Approximate Solution of $RNDP_{UO}(D)$ Based on the Price of Anarchy

In light of the comparative difficulty of solving $RNDP_{UO}(D)$ to solving $RNDP_{SO}(D)$, we establish in this section that the optimal solution to the latter problem offers an approximate solution to the former. The accuracy of this approximation is given by the price of anarchy. In this way, we extend the NDP approximation result of Roughgarden [140] to the context of NDPs involving nonseparable arc costs and to the robust NDP setting.

5.4.1 Price of Anarchy Review

Let us begin by briefly reviewing some definitions and results from the literature on the price of anarchy. This measure of inefficiency is defined as the ratio of the highest total system cost under any equilibrium to the minimal total system cost, i.e., the cost under a SO solution. In the following definition, we specialize the description of the price of anarchy to the NDP context. That is, this definition takes into account the dependence of UO and SO solutions on the arc design vector $\mathbf{y}$ and on the vector $\mathbf{d}$ of demands for the O-D pairs.

Definition 5.4.1. Consider a network $G(V, A)$ with vector of arc cost functions $\mathbf{c}$. For a given $\mathbf{y} \in Y$, the **price of anarchy** of the resulting network, under given demand rates $\mathbf{d} \in \mathbb{R}_+^{|W|}$, is defined as

$$\frac{\zeta_{UO}(\mathbf{y}, \mathbf{d})}{\zeta_{SO}(\mathbf{y}, \mathbf{d})},$$

where $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ and $\zeta_{SO}(\mathbf{y}, \mathbf{d})$ are defined as in Section 5.2.2.

As mentioned in Section 5.1.2, several contributors to the literature have derived tight bounds on the price of anarchy that are independent of network topology and of $\mathbf{d}$. For example, Perakis [127] derived a tight bound, whose exact value depends on the degree of nonlinearity, denoted by κ, and the degree of asymmetry, denoted by γ^2, of the vector of arc cost functions.[2] Specifically, Perakis considered the more general setting of $VI(K, \mathbf{c})$, with $K \subseteq \mathbb{R}^n$. She defined

$$\gamma^2 \;=\; \sup_{\mathbf{x} \in K} ||\mathbf{S}(\mathbf{x})^{-1} J\mathbf{c}(\mathbf{x})||^2_{\mathbf{S}(\mathbf{x})},$$

where

$$\mathbf{S}(\mathbf{x}) \;=\; \frac{J\mathbf{c}(\mathbf{x}) + [J\mathbf{c}(\mathbf{x})]'}{2}.$$

[2]Perakis [127] denotes the degrees of nonlinearity and asymmetry by A and c^2, respectively. In this chapter of the thesis, we use A to refer to the arc set of a graph and $\mathbf{c}$ to denote the vector of cost functions. Therefore, in order to avoid confusion, we use κ and γ^2 to discuss the measures of nonlinearity and asymmetry.

In addition, she let κ be a constant greater than or equal to one and satisfying

$$\frac{1}{\kappa}\boldsymbol{\mu}' J\mathbf{c}(\mathbf{x})\boldsymbol{\mu} \;\leq\; \boldsymbol{\mu}' J\mathbf{c}(\mathbf{z})\boldsymbol{\mu} \;\leq\; \kappa\boldsymbol{\mu}' J\mathbf{c}(\mathbf{x})\boldsymbol{\mu}, \qquad \forall \mathbf{x}, \mathbf{z} \in K;\ \forall \boldsymbol{\mu} \in \mathbb{R}^n.$$

Perakis proved that, in this general $VI(K, \mathbf{c})$ setting, the price of anarchy of a network can be no greater than

$$\alpha \;=\; \begin{cases} \frac{4}{4-\gamma^2\kappa}, & \gamma^2 \leq 2/\kappa, \\ \gamma^2\kappa^2 - 2(\kappa - 1), & \gamma^2 > 2/\kappa. \end{cases}$$

In the context of network flows, it is easy to show that, if the arc cost functions are affine, then $\kappa = 1$. Furthermore, if the arc cost functions are symmetric, $\gamma^2 = 1$. As a result, in the case of affine and symmetric arc costs, Perakis' bound simplifies to the bound of 4/3 due to Roughgarden and Tardos [143, 144] and Roughgarden [141].

Consider the context of a network whose arc set $A(\mathbf{y})$ is determined by a vector, $\mathbf{y} \in Y$, of binary design variables. In this setting, $\mathbf{y}$ thereby also determines the collection of arc cost functions to be associated with the resulting network and, consequently, the corresponding degrees of nonlinearity and asymmetry, $\kappa(\mathbf{y})$ and $\gamma^2(\mathbf{y})$, respectively. As a result, for networks whose arc set is a function of $\mathbf{y}$, Perakis' bound is also a function of $\mathbf{y}$. We denote this $\mathbf{y}$-dependent version of Perakis' bound by $\alpha(\mathbf{y})$. To clarify, this bound depends on $\mathbf{y}$ only because $\mathbf{y}$ determines the collection of, and therefore the degrees of nonlinearity and asymmetry of, the arc cost functions. The dependence is *not* because $\mathbf{y}$ determines the network topology as well. In some settings, $\alpha(\mathbf{y})$ is a constant $\forall \mathbf{y} \in Y$. For instance, $\kappa(\mathbf{e}) = 1$, where $\mathbf{e} \in \mathbb{R}^{|A|}$ denotes the vector of all ones, implies $\kappa(\mathbf{y}) = 1$, $\forall \mathbf{y} \in \{0,1\}^{|A|}$. Similarly, $\gamma^2(\mathbf{e}) = 1$ implies $\gamma^2(\mathbf{y}) = 1$, $\forall \mathbf{y} \in \{0,1\}^{|A|}$. Accordingly, in settings with $\kappa(\mathbf{e}) = \gamma^2(\mathbf{e}) = 1$, $\alpha(\mathbf{y}) = 4/3$, $\forall \mathbf{y} \in Y$.

As discussed in Section 5.1.3, Roughgarden [140] considered a version of the nominal NDP involving no construction budget constraint and separable arc cost functions. In terms of the notation of Section 5.2.2, he considered $NDP_{UO}(\check{\mathbf{d}})$ with $B \geq \sum_{a \in A} b_a$, and $J\mathbf{c}(\mathbf{f})$ a diagonal matrix. For $G(V, A)$ and W given, he established the following result, which we rephrase in terms of our notation. Let $\alpha \geq 1$ be a bound on the price of anarchy for $G(V, A(\mathbf{y})), \forall \mathbf{y} \in Y$. Then, setting $\mathbf{y} = \mathbf{e}$, where $\mathbf{e} \in \mathbb{R}^{|A|}$ is the vector of all ones (denoting that all arcs be built), is an α-approximation algorithm for the aforementioned "unbudgeted" NDP with separable arc costs. To address cases with prohibitively high α, Roughgarden proved an alternate approximation bound of $\frac{|V|}{2}$, where V is the set of nodes in the network. He showed that building the entire arc set is also a $\frac{|V|}{2}$-approximation algorithm for the NDP, and that no better approximation algorithm exists, unless $P = NP$. In essence, Roughgarden established that, in the absence of a budget constraint, one cannot efficiently find a better approximate solution to $NDP_{UO}(\check{\mathbf{d}})$ than the solution of building the entire arc set.

5.4.2 An Approximation Result

In the nominal or robust NDP with active construction budget constraint, i.e., with $B < \sum_{a \in A} b_a$, $\mathbf{y} = \mathbf{e}$ may be an infeasible solution, if it is too costly to build the entire edge set. In the following theorem, we extend Roughgarden's approximation result to the NDP with nonseparable arc costs and with active construction budget constraint. This theorem also generalizes Roughgarden's result to the robust NDP setting.

Theorem 5.4.1. *Consider a network $G(V, A)$, with O-D pair set W and a vector of continuous arc cost functions $\mathbf{c}$. Suppose that $D \subseteq \mathbb{R}_+^{|W|}$, the uncertainty set of demands $\tilde{\mathbf{d}}$, is closed and bounded. Let $\alpha(\mathbf{y})$ denote an upper bound on the price of anarchy of subnetwork $G(V, A(\mathbf{y}))$ and*

$$\begin{aligned} \alpha = \max_{\mathbf{y}} \quad & \alpha(\mathbf{y}) \\ s.t. \quad & \mathbf{y} \in Y. \end{aligned}$$

Furthermore, let $\mathbf{y}_{SO}^*$ and $\mathbf{y}_{UO}^*$ denote optimal solutions of $RNDP_{SO}(D)$ and $RNDP_{UO}(D)$, *respectively. That is,*

$$\begin{aligned}\mathbf{y}^*_{SO} &\in \arg\min_{\mathbf{y}} \left\{ \tau_{SO}(D;\mathbf{y}) \;\middle|\; \mathbf{y} \in \{0,1\}^{|A|},\ \mathbf{b}'\mathbf{y} \le B,\ K_A(\mathbf{y},\mathbf{e}) \ne \emptyset \right\},\\ \mathbf{y}^*_{UO} &\in \arg\min_{\mathbf{y}} \left\{ \tau_{UO}(D;\mathbf{y}) \;\middle|\; \mathbf{y} \in \{0,1\}^{|A|},\ \mathbf{b}'\mathbf{y} \le B,\ K_A(\mathbf{y},\mathbf{e}) \ne \emptyset \right\}.\end{aligned}$$

Then, $\mathbf{y}^*_{SO}$ is α-optimal for $RNDP_{UO}(D)$, i.e.,

$$\tau_{UO}(D;\mathbf{y}^*_{UO}) \le \tau_{UO}(D;\mathbf{y}^*_{SO}) \le \alpha \cdot \tau_{UO}(D;\mathbf{y}^*_{UO}). \tag{5.20}$$

Proof. The first inequality (5.20) holds since $\mathbf{y}^*_{SO}$ is feasible, but not necessarily optimal for $RNDP_{UO}(D)$. $\forall \mathbf{y} \in Y$, let

$$\begin{aligned}\mathbf{d}_{SO}(\mathbf{y}) &\in \arg\max_{\mathbf{d}\in D} \zeta_{SO}(\mathbf{y},\mathbf{d}),\\ \mathbf{d}_{UO}(\mathbf{y}) &\in \arg\max_{\mathbf{d}\in D} \zeta_{UO}(\mathbf{y},\mathbf{d}),\\ \mathbf{f}_{SO}(\mathbf{y},\mathbf{d}) &\in \arg\min_{\mathbf{f}} \left\{ \mathbf{c}(\mathbf{f})'\mathbf{f} \;\middle|\; \mathbf{f} \in K_A(\mathbf{y},\mathbf{d}) \right\} = SO_A(\mathbf{y},\mathbf{d}),\\ \mathbf{f}_{UO}(\mathbf{y},\mathbf{d}) &\in \arg\max_{\mathbf{f}} \left\{ \mathbf{c}(\mathbf{f})'\mathbf{f} \;\middle|\; \mathbf{f} \in UO_A(\mathbf{y},\mathbf{d}) \right\}.\end{aligned} \tag{5.21}$$

In other words, $\mathbf{d}_{SO}(\mathbf{y})$ and $\mathbf{d}_{UO}(\mathbf{y})$ denote the vectors of demands giving rise to the worst-case total system costs under SO and UO routing, respectively. For a given $\mathbf{y} \in Y$ and a given $\mathbf{d}$, $\mathbf{f}_{SO}(\mathbf{y},\mathbf{d})$ denotes a corresponding SO flow solution, and $\mathbf{f}_{UO}(\mathbf{y},\mathbf{d})$ is a corresponding UO flow solution giving rise to the highest possible total system cost.

$\forall \mathbf{y} \in Y$,

$$\begin{aligned}\tau_{UO}(D;\mathbf{y}) &= \mathbf{c}(\mathbf{f}_{UO}(\mathbf{y},\mathbf{d}_{UO}))'\,\mathbf{f}_{UO}(\mathbf{y},\mathbf{d}_{UO})\\ &\le \alpha \cdot \mathbf{c}(\mathbf{f}_{SO}(\mathbf{y},\mathbf{d}_{UO}))'\,\mathbf{f}_{SO}(\mathbf{y},\mathbf{d}_{UO})\\ &\le \alpha \cdot \mathbf{c}(\mathbf{f}_{SO}(\mathbf{y},\mathbf{d}_{SO}))'\,\mathbf{f}_{SO}(\mathbf{y},\mathbf{d}_{SO})\\ &= \alpha \cdot \tau_{SO}(D;\mathbf{y}).\end{aligned}$$

The equalities hold by the definitions of $\tau_{UO}(D;\mathbf{y})$ and $\tau_{SO}(D;\mathbf{y})$, respectively. The first inequality follows from the definition of α, and the second inequality holds since $\alpha \geq 0$ and since $\mathbf{d}_{UO}(\mathbf{y}) \in D$ but need not belong to

$$\arg\max_{\mathbf{d}\in D} \zeta_{SO}(\mathbf{y},\mathbf{d}).$$

Therefore, taking $\mathbf{y} = \mathbf{y}^*_{SO}$, we obtain

$$\tau_{UO}\left(D;\mathbf{y}^*_{SO}\right) \;\leq\; \alpha \cdot \tau_{SO}\left(D;\mathbf{y}^*_{SO}\right).$$

To complete the proof, we need to show that $\tau_{SO}\left(D;\mathbf{y}^*_{SO}\right) \leq \tau_{UO}\left(D;\mathbf{y}^*_{UO}\right)$. To begin, $\forall \mathbf{y} \in \{0,1\}^{|A|}$, and $\forall \mathbf{d} \in \mathbb{R}^{|W|}_{+}$,

$$UO_A(\mathbf{y},\mathbf{d}) \;\subseteq\; K_A(\mathbf{y},\mathbf{d}),$$

but it may be the case that

$$UO_A(\mathbf{y},\mathbf{d}) \cap SO_A(\mathbf{y},\mathbf{d}) \;=\; \emptyset.$$

Therefore,

$$\zeta_{SO}(\mathbf{y},\mathbf{d}) \;\leq\; \zeta_{UO}(\mathbf{y},\mathbf{d}).$$

Taking the maximum over $\mathbf{d} \in D$ of both sides, we obtain

$$\tau_{SO}(D;\mathbf{y}) \;\leq\; \tau_{UO}(D;\mathbf{y}).$$

Taking the minimum over $\mathbf{y} \in \{0,1\}^{|A|}$ such that $\mathbf{b}'\mathbf{y} \leq B$ and $K_A(\mathbf{y},\mathbf{e}) \neq \emptyset$, the desired result follows. □

Remark: As made clear in the proof, Theorem 5.4.1 is not an immediate consequence of the definition of the price of anarchy. Indeed, the price of anarchy compares costs induced by a *single design vector* $\mathbf{y} \in Y$, under *UO versus SO routing.* That is, as an immediate consequence of the definition,

$$\tau_{SO}(D;\mathbf{y}) \quad \leq \quad \tau_{UO}(D;\mathbf{y}) \leq \alpha \cdot \tau_{SO}(D;\mathbf{y}).$$

In contrast, Theorem 5.4.1 compares the costs induced by *two different design vectors,* $\mathbf{y}^*_{UO}$ versus $\mathbf{y}^*_{SO}$, *both under UO routing.* For the purpose of a side-by-side comparison, recall that the theorem establishes that

$$\tau_{UO}(D;\mathbf{y}^*_{UO}) \quad \leq \quad \tau_{UO}(D;\mathbf{y}^*_{SO}) \leq \alpha \cdot \tau_{UO}(D;\mathbf{y}^*_{UO}).$$

In addition, note that, when $D = \{\check{\mathbf{d}}\}$, i.e., is a singleton, the theorem says that the optimal solution of $NDP_{SO}(\check{\mathbf{d}})$ is α-optimal for $NDP_{UO}(\check{\mathbf{d}})$. Furthermore, $\forall \mathbf{y} \in \{0,1\}^{|A|}$,

$$K_A(\mathbf{y},\mathbf{d}) \quad \subseteq \quad K_A(\mathbf{e},\mathbf{d}),$$

where $\mathbf{e} \in \mathbb{R}^{|A|}$ denotes the vector of all ones. Accordingly, in the context of the unbudgeted version of the nominal NDP, $\mathbf{y}^*_{SO} = \mathbf{e}$, as in Roughgarden's approximation result under separable arc costs.

In the NDP setting, either robust or nominal, when there is no construction budget, it requires virtually no computational effort to set $\mathbf{y}^*_{SO} = \mathbf{e}$, the vector of all ones. In contrast, when $B < \sum_{a \in A} b_a$, and $\mathbf{y}$ must satisfy the budget constraint $\mathbf{b}'\mathbf{y} \leq B$, computing $\mathbf{y}^*_{SO}$ is, in general, quite difficult. Thus, while Theorem 5.4.1 identifies an approximate solution of $RNDP_{UO}(D)$, it does not identify an approximation algorithm. Indeed, $\mathbf{y}^*_{SO}$ represents the optimal solution to an optimization problem with binary integer constraints. Furthermore, in the robust setting, as we discussed in Section 5.3, simply evaluating $\tau_{SO}(D;\mathbf{y})$ requires the solution of a bilevel optimization problem, which is, in general, NP-hard.

Nonetheless, as the discussion in Section 5.3 also suggested, evaluating $\tau_{SO}(D;\mathbf{y})$ should generally, from a practical standpoint, be significantly easier than computing $\tau_{UO}(D;\mathbf{y})$. Accordingly, solving $RNDP_{SO}(D)$ and thereby identifying $\mathbf{y}^*_{SO}$ may be significantly easier than solving $RNDP_{UO}(D)$. Thus, the approximation result presented in Theorem 5.4.1 may prove useful, especially in cases in which α is relatively close to one.

5.5 A Branch and Bound Algorithm for the Robust NDP

Recall the two fundamental requirements for solving any optimization problem. At a high level, one must navigate through the feasible region in order to identify an optimal solution. At a lower level, one must evaluate the performance of given feasible solutions visited along the way. In the preceding parts of this chapter, we have discussed methods for carrying out only the latter task. Indeed, in Section 5.3, we proposed ways of computing $\tau_{SO}(D;\mathbf{y})$ and $\tau_{UO}(D;\mathbf{y})$, the worst-case performances of design decision vector $\mathbf{y}$ under SO and UO routing, respectively. In this section, we address the high level task of navigating through the feasible region of the robust NDP in search of an optimal solution.

Specifically, we propose a branch and bound algorithm for solving $RNDP_{SO}(D)$ and $RNDP_{UO}(D)$. When $\tau_{SO}(D;\mathbf{y})$ and $\tau_{UO}(D;\mathbf{y})$ can be computed exactly, the algorithm determines a connectivity and budget feasible design, $\mathbf{y}^*$, minimizing $\tau_{SO}(D;\mathbf{y})$ for $RNDP_{SO}(D)$ and $\tau_{UO}(D;\mathbf{y})$ for $RNDP_{UO}(D)$. It outputs this optimal design solution and the corresponding optimal objective value, either $\tau_{SO}(D;\mathbf{y}^*)$ or $\tau_{UO}(D;\mathbf{y}^*)$. For cases in which $\tau_{UO}(D;\mathbf{y})$ cannot be computed exactly but can be upper-bounded, we propose a variation of the algorithm that heuristically solves $RNDP_{UO}(D)$. In particular, this heuristic version of the algorithm determines a feasible design solution $\mathbf{y}^*$ that minimizes this upper bound. In this setting, the method outputs the resulting heuristic solution $\mathbf{y}^*$ and a range to which the corresponding exact value of $\tau_{UO}(D;\mathbf{y}^*)$, though it cannot be computed precisely, is guaranteed to belong.

Our branch and bound method is based on LeBlanc's [87] branch and bound algorithm for the nominal NDP under UO routing. However, our approach differs from LeBlanc's work in several ways. First, and most importantly, as already noted, LeBlanc's method applies only to the nominal NDP under UO routing. In contrast, our algorithm and our analysis in this section address the *robust* NDP under either SO or UO routing. Therefore, our method covers the nominal NDP under UO routing, but as a special case of the more general robust setting. Second, while LeBlanc's method uses only lower bounds, we not only use analogous lower bounds but also suggest upper bounds that may be helpful in pruning the branch and bound tree. In particular, in the context of UO routing, our upper bounds are based on the price of anarchy. In this way, our algorithm represents the first prescriptive, rather than descriptive, use of the concept of the price of anarchy. Third, in presenting his treatment of the NDP, LeBlanc discussed only problems involving vectors of separable arc-usage cost functions. In fact, his method also applies to NDPs involving nonseparable cost functions. The discussion we give here directly addresses the general case of nonseparable and possibly even asymmetric arc cost functions.

5.5.1 Notation and Definitions

Before presenting our branch and bound algorithm for the robust NDP, let us establish some notation and definitions. For an introduction to branch and bound algorithms for optimization problems involving integer constraints, we refer the interested reader to Bertsimas and Weismantel [20].

The algorithm we present in this section induces a branch and bound tree. Nodes in the tree represent varying levels of commitment with respect to design decisions. For instance, each leaf node in the tree corresponds to a fully-specified design vector $\mathbf{y} \in \{0,1\}^{|A|}$. Recall that $y_a = 1$ denotes that arc $a \in A$ is to be constructed, and $y_a = 0$ denotes that it is not. In contrast, each non-leaf node corresponds to a partially specified design vector $\boldsymbol{\pi} \in \{-1,0,1\}^{|A|}$. As with fully-specified design vectors, $\pi_a = 1$ and $\pi_a = 0$ denote, respectively, that arc $a \in A$ is or is not to be constructed. In addition, $\pi_a = -1$ denotes that arc $a \in A$

may or may not be constructed. For the sake of brevity, we will refer to fully-specified and partially-specified design solutions as full and partial solutions, respectively.

Definition 5.5.1. Let $S(\boldsymbol{\pi})$ denote the set of full solutions that are **successors** or **descendants** of $\boldsymbol{\pi}$. That is,

$$S(\boldsymbol{\pi}) = \left\{ \mathbf{y} \in \{0,1\}^{|A|} \,\middle|\, y_a = \pi_a, \ \forall a \in A \text{ s.t. } \pi_a \in \{0,1\} \right\}.$$

Furthermore, define the **yes-completion** of $\boldsymbol{\pi}$ to be the unique $\mathbf{y} \in \{0,1\}^{|A|}$ such that, $\forall a \in A$,

$$y_a = \begin{cases} \pi_a, & \pi_a \in \{0,1\}, \\ 1, & \pi_a = -1. \end{cases}$$

Similarly, define the **no-completion** of $\boldsymbol{\pi}$ to be the unique $\mathbf{y} \in \{0,1\}^{|A|}$ such that, $\forall a \in A$,

$$y_a = \begin{cases} \pi_a, & \pi_a \in \{0,1\}, \\ 0, & \pi_a = -1. \end{cases}$$

Figure 5-4 illustrates an example, in which $|A| = 5$, of a subtree of a branch and bound tree. This subtree is rooted at $\boldsymbol{\pi} = (-1, 1, -1, -1, 0)$ and contains all of the corresponding successor partial and full solutions, including the yes- and no-completions of $\boldsymbol{\pi}$.

Recall that not all design vectors $\mathbf{y} \in \{0,1\}^{|A|}$, and therefore not all leaf nodes in the branch and bound tree, are necessarily feasible with respect to the construction budget and the connectivity of the O-D pairs. As before, let Y denote the set of connectivity feasible full solutions, i.e.,

$$Y = \left\{ \mathbf{y} \in \{0,1\}^{|A|} \,\middle|\, K_A(\mathbf{y}, \mathbf{e}) \neq \emptyset \right\},$$

where $\mathbf{e} \in \mathbb{R}^{|W|}$ denotes the vector of all ones.

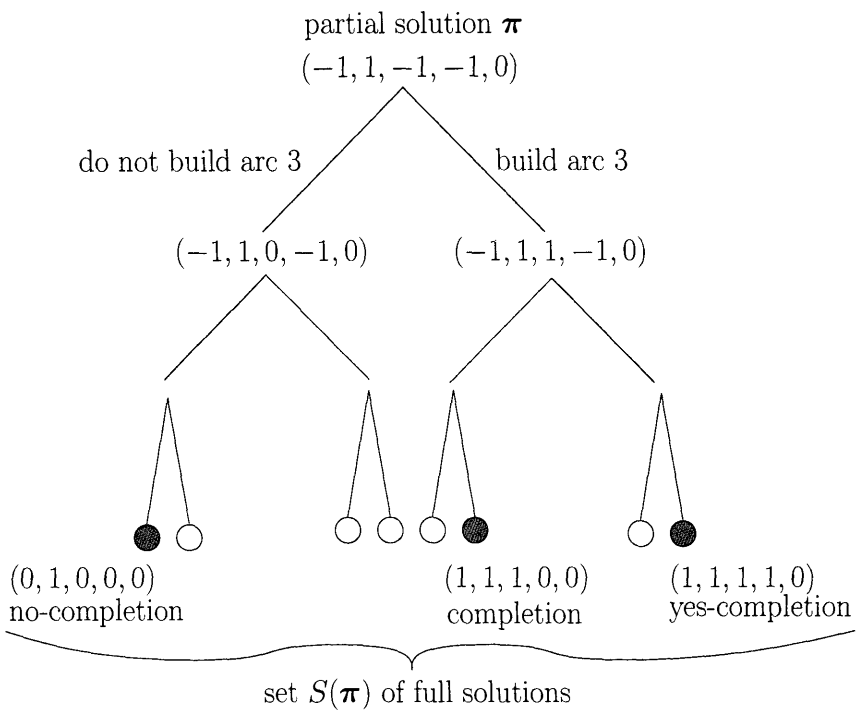

Figure 5-4: A branch and bound subtree rooted at $\boldsymbol{\pi}$

5.5.2 Bounding the Performance of a Partial Solution's Descendants

Having set forth the notation and terminology we will use in presenting our branch and bound algorithm, we now establish the bounds we will use in the algorithm.

First, consider the context of SO routing. For any partial design solution $\boldsymbol{\pi} \in \{-1,0,1\}^{|A|}$, we wish to lower- and upper-bound the worst-case performance of any descendant of $\boldsymbol{\pi}$. As one would intuitively expect, $\forall \boldsymbol{\pi} \in \{-1,0,1\}^{|A|}$, the best- and worst-performing descendants of $\boldsymbol{\pi}$, under SO routing, are its yes- and no-completions, respectively.

Lemma 5.5.1. *$\forall \boldsymbol{\pi} \in \{-1,0,1\}^{|A|}$, let $\mathbf{y}^1$ and $\mathbf{y}^0$ denote the yes- and no-completions, respectively, of $\boldsymbol{\pi}$. Let $D \subseteq \mathbb{R}_+^{|W|}$. If $\mathbf{y}^1 \in Y$, then*

$$\tau_{SO}\left(D;\mathbf{y}^1\right) \;\leq\; \tau_{SO}\left(D;\mathbf{y}\right), \qquad \forall \mathbf{y} \in S(\boldsymbol{\pi}) \cap Y. \tag{5.22}$$

In addition, if $\mathbf{y}^0 \in Y$*, then*

$$\tau_{SO}(D;\mathbf{y}) \;\leq\; \tau_{SO}\left(D;\mathbf{y}^0\right), \qquad \forall \mathbf{y} \in S(\boldsymbol{\pi}) \cap Y. \tag{5.23}$$

Proof. Suppose $\mathbf{y}^1 \in Y$ and consider any $\mathbf{y} \in S(\boldsymbol{\pi}) \cap Y$. $\forall \mathbf{d} \in \mathbb{R}_+^{|W|}$, $K_A(\mathbf{y}, \mathbf{d}) \subseteq K_A(\mathbf{y}^1, \mathbf{d})$, thereby implying (5.22). Suppose, in addition, that $\mathbf{y}^0 \in Y$. $\forall \mathbf{d} \in \mathbb{R}_+^{|W|}$, $K_A(\mathbf{y}^0, \mathbf{d}) \subseteq K_A(\mathbf{y}, \mathbf{d})$, thereby implying (5.23). □

Next, consider the context of UO routing. Recall Braess' Paradox [27], as discussed in Section 5.1.3. Braess' example demonstrates that, in the nominal NDP setting, any partial design solution $\boldsymbol{\pi}$ may have a completion that performs strictly better than the corresponding yes-completion. Since the nominal NDP is a special instance of the robust NDP, this result extends to the robust NDP context as well. Accordingly, for a given $\boldsymbol{\pi} \in \{-1, 0, 1\}^{|A|}$ for which $S(\boldsymbol{\pi}) \cap Y \neq \emptyset$, there may be a $\mathbf{y} \in S(\boldsymbol{\pi}) \cap Y$ such that $\tau_{UO}(D;\mathbf{y}^1) > \tau_{UO}(D;\mathbf{y})$, where $\mathbf{y}^1$ is the yes-completion of $\boldsymbol{\pi}$. For this reason, analogous to LeBlanc's [87] lower bound in the nominal NDP setting, we lower-bound $\tau_{UO}(D;\mathbf{y})$ by $\tau_{SO}(D;\mathbf{y}^1)$. Unlike LeBlanc, we additionally prove an upper bound on $\tau_{UO}(D;\mathbf{y})$. Like the aforementioned lower bound, this upper bound is also derived from a comparison of performance under UO and SO routing.

Lemma 5.5.2. *$\forall \boldsymbol{\pi} \in \{-1, 0, 1\}^{|A|}$, let $\mathbf{y}^1$ and $\mathbf{y}^0$ denote the yes- and no-completions, respectively, of $\boldsymbol{\pi}$. Let $D \subseteq \mathbb{R}_+^{|W|}$. If $\mathbf{y}^1 \in Y$, then*

$$\tau_{SO}\left(D;\mathbf{y}^1\right) \;\leq\; \tau_{UO}(D;\mathbf{y}), \qquad \forall \mathbf{y} \in S(\boldsymbol{\pi}) \cap Y.$$

Let α be a bound on the price of anarchy for $G(V, A(\mathbf{y}))$, $\forall \mathbf{y} \in Y$. If $\mathbf{y}^0 \in Y$, then

$$\tau_{UO}(D;\mathbf{y}) \;\leq\; \alpha \cdot \tau_{SO}\left(D;\mathbf{y}^0\right), \qquad \forall \mathbf{y} \in S(\boldsymbol{\pi}) \cap Y.$$

Proof. $\forall \mathbf{y} \in Y$, let $\mathbf{d}_{SO}(\mathbf{y})$ and $\mathbf{d}_{UO}(\mathbf{y})$ be defined as in (5.21). In addition, $\forall \mathbf{y} \in Y$, $\forall \mathbf{d} \in \mathbb{R}_+^{|W|}$, let $\mathbf{f}_{SO}(\mathbf{y}, \mathbf{d})$ and $\mathbf{f}_{UO}(\mathbf{y}, \mathbf{d})$ be defined as in (5.21). Suppose $\mathbf{y}^1 \in Y$ and consider any $\mathbf{y} \in S(\boldsymbol{\pi}) \cap Y$. We already showed, in the course of the proof of Theorem 5.4.1, that $\forall \mathbf{y} \in Y$,

$$\tau_{SO}(D; \mathbf{y}) \;\leq\; \tau_{UO}(D; \mathbf{y}) \;\leq\; \alpha \cdot \tau_{SO}(D; \mathbf{y}).$$

Therefore, the desired result follows from Lemma 5.5.1. □

5.5.3 The Branch and Bound Algorithm

We now present the branch and bound algorithm for $RNDP_{SO}(D)$ and $RNDP_{UO}(D)$. So that we may cover both problems simultaneously, let us introduce some additional notation. Let $\rho \in \{SO, UO\}$ denote a routing paradigm, either system-optimal or user-optimal. For $\boldsymbol{\pi} \in \{-1, 0, 1\}$, such that $S(\boldsymbol{\pi}) \cap Y \neq \emptyset$, let $LB_\rho(S(\boldsymbol{\pi}))$ denote a lower bound on $\tau_\rho(D; \mathbf{y})$, $\forall \mathbf{y} \in S(\boldsymbol{\pi}) \cap Y$. In particular, in accordance with Lemmas 5.5.1 and 5.5.2, let

$$LB_\rho(S(\boldsymbol{\pi})) \;=\; \tau_{SO}\left(D; \mathbf{y}^1\right), \qquad \rho \in \{SO, UO\},$$

where $\mathbf{y}^1$ is the yes-completion of $\boldsymbol{\pi}$. Similarly, if $\mathbf{y}^0$ is the no-completion of $\boldsymbol{\pi}$ and $\mathbf{y}^0 \in Y$, then let $UB_\rho(S(\boldsymbol{\pi}))$ denote an upper bound on $\tau_\rho(D; \mathbf{y})$, $\forall \mathbf{y} \in S(\boldsymbol{\pi}) \cap Y$. Specifically, as suggested by Lemmas 5.5.1 and 5.5.2, let

$$UB_\rho(S(\boldsymbol{\pi})) \;=\; \begin{cases} \tau_{SO}\left(D; \mathbf{y}^0\right), & \rho = SO, \\ \alpha \cdot \tau_{SO}\left(D; \mathbf{y}^0\right), & \rho = UO. \end{cases}$$

For clarity and conciseness of exposition, in presenting the steps of our branch and bound algorithm, we initially assume that $\tau_{SO}(D; \mathbf{y})$ and $\tau_{UO}(D; \mathbf{y})$ can both be computed exactly. Under these conditions, for $RNDP_{SO}(D)$ or $RNDP_{UO}(D)$, the algorithm determines the connectivity and budget feasible design solution, $\mathbf{y}^*$, that exactly minimizes $\tau_{SO}(D; \mathbf{y})$ or

$\tau_{UO}(D;\mathbf{y})$, respectively. It outputs this optimal solution and the corresponding objective value.

After presenting this exact version of the algorithm, we describe alterations to the algorithm that enable the heuristic solution of $RNDP_{UO}(D)$, when $\tau_{UO}(D;\mathbf{y})$ cannot be computed exactly but can be upper-bounded. Under these conditions, the altered algorithm determines the connectivity and budget feasible design solution, $\mathbf{y}^*$, that minimizes the aforementioned upper bound on $\tau_{UO}(D;\mathbf{y})$. It outputs both this heuristic solution, $\mathbf{y}^*$, and a range of values to which the corresponding $\tau_{UO}(D;\mathbf{y}^*)$ is guaranteed to belong.

Robust NDP Branch and Bound Algorithm.

Input: node set V

virtual arc set A

set W of O-D pairs

demand uncertainty set D

vector function $\mathbf{c}:\mathbb{R}^{|A|}\to\mathbb{R}^{|A|}$ giving the costs per unit flow

vector of design costs $\mathbf{b}$

design budget B

a bound α on the price of anarchy for any subnetwork of $G(V,A)$

Output: a connectivity and budget feasible $\mathbf{y}$ minimizing $\tau_\rho(D;\mathbf{y})$

the corresponding value of $\tau_\rho(D;\mathbf{y})$

Step 1. Initialize.

a. Set $\Omega := \{-\mathbf{e}\}$. where $\mathbf{e}\in\mathbb{R}^{|A|}$ is the vector of all ones.

b. Set $UB^* := \infty$.

c. For any $\mathbf{y}\in Y$ such that $\mathbf{b}'\mathbf{y}\le B$, set

$$\begin{aligned}\mathbf{y}^* &:= \mathbf{y},\\ OPT &:= \tau_\rho(D;\mathbf{y}).\end{aligned}$$

Step 2. If $\Omega=\emptyset$, then output $\mathbf{y}^*$, output OPT, and stop.

Otherwise, branch as follows.

a. Select $\boldsymbol{\pi} \in \Omega$ and $a' \in A$ such that $\pi_{a'} = -1$. Set

$$\Omega \;:=\; \Omega \backslash \{\boldsymbol{\pi}\}.$$

If $LB_\rho\left(S(\boldsymbol{\pi})\right) \geq \min\{OPT, UB^*\}$, then restart step 2. (All descendants of $\boldsymbol{\pi}$ are no better than the current incumbent solution $\mathbf{y}^*$ or the optimal descendant of the partial solutions in Ω).

b. Set $\boldsymbol{\pi}^i$, $i \in \{0,1\}$ according to

$$\pi_a^i \;:=\; \begin{cases} i, & a = a', \\ \pi_a, & a \neq a'. \end{cases} \tag{5.24}$$

Step 3. Consider $\boldsymbol{\pi}^0$:

a. Let $\mathbf{y}^{0,1}$ denote the yes-completion of $\boldsymbol{\pi}^0$. If $K_A\left(\mathbf{y}^{0,1}, \mathbf{e}\right) = \emptyset$, where $\mathbf{e} \in \mathbb{R}^{|W|}$ is the vector of all ones, then proceed to step 4. ($S\left(\boldsymbol{\pi}^0\right) \cap Y = \emptyset$).

b. If $\{a \in A \mid \pi_a^0 = -1\} = \emptyset$ (i.e., $\boldsymbol{\pi}^0$ is a full solution), and if $\tau_\rho\left(D; \boldsymbol{\pi}^0\right) < OPT$, then set

$$\begin{aligned} \mathbf{y}^* \;&:=\; \boldsymbol{\pi}^0, \\ OPT \;&:=\; \tau_\rho\left(D; \boldsymbol{\pi}^0\right). \end{aligned}$$

c. If $\{a \in A \mid \pi_a^0 = -1\} \neq \emptyset$ (i.e., $\boldsymbol{\pi}^0$ is not a full solution), and if $LB_\rho\left(S\left(\boldsymbol{\pi}^0\right)\right) < \min\{OPT, UB^*\}$, then set $\Omega := \Omega \cup \{\boldsymbol{\pi}^0\}$.

Step 4. Consider $\boldsymbol{\pi}^1$:

a. Let $\mathbf{y}^{1,0}$ denote the no-completion of $\boldsymbol{\pi}^1$. If $\mathbf{b}'\left(\mathbf{y}^{1,0}\right) > B$, then go to step 2. (None of the descendants of $\boldsymbol{\pi}^1$ is budget-feasible).

b. If $\{a \in A \mid \pi_a^1 = -1\} = \emptyset$ (i.e., $\boldsymbol{\pi}^1$ is a full solution), and if $\tau_\rho(D; \boldsymbol{\pi}^1) < OPT$, then set

$$\begin{aligned} \mathbf{y}^* &:= \boldsymbol{\pi}^1 \\ OPT &:= \tau_\rho\left(D; \boldsymbol{\pi}^1\right). \end{aligned}$$

c. If $\{a \in A \mid \pi_a^1 = -1\} \neq \emptyset$ (i.e., $\boldsymbol{\pi}^1$ is not a full solution), and if $LB_\rho(S(\boldsymbol{\pi}^1)) = LB_\rho(S(\boldsymbol{\pi})) < \min\{OPT, UB^*\}$, then set $\Omega := \Omega \cup \{\boldsymbol{\pi}^1\}$.

d. If $\mathbf{y}^{1,0} \in Y$, then set $UB^* := \min\{UB^*,\ UB_\rho(S(\boldsymbol{\pi}^1))\}$.

Step 5. Go to step 2.

From Lemmas 5.5.1 and 5.5.2, the correctness of this branch and bound algorithm, when $\tau_{SO}(D; \mathbf{y})$ and $\tau_{UO}(D; \mathbf{y})$ can be computed exactly, immediately follows.

Theorem 5.5.1. *The Robust NDP Branch and Bound Algorithm correctly solves $RNDP_\rho(D)$, $\forall \rho \in \{SO, UO\}$.*

Example 5.5.1. To better illustrate the features of and ideas behind our branch and bound algorithm, consider the following simple example of $RNDP_{UO}(D)$, pictured in Figure 5-5. In this example, there are two disjoint O-D pairs. The first is connected by arcs 1 (top) and 2 (bottom), and the second is connected by arcs 3 (top) and 4 (bottom). The vector of arc construction costs is given by $\mathbf{b} = (1, 2, 1, 2)$. The arc cost functions are given by

$$c_a(\mathbf{f}) = \begin{cases} 100, & a \in \{1, 3\} \\ 10 f_a, & a \in \{2, 4\}. \end{cases}$$

Suppose the construction budget $B = 4$ and the demand uncertainty set D is given by

$$D = \left\{ \tilde{\mathbf{d}} \in \mathbb{R}^{|W|} \;\middle|\; \tilde{d}_1 \in [8, 12],\ \tilde{d}_2 \in [9, 15];\ \frac{\tilde{d}_1 - 8}{4} + \frac{\tilde{d}_2 - 9}{6} \leq 1 \right\}.$$

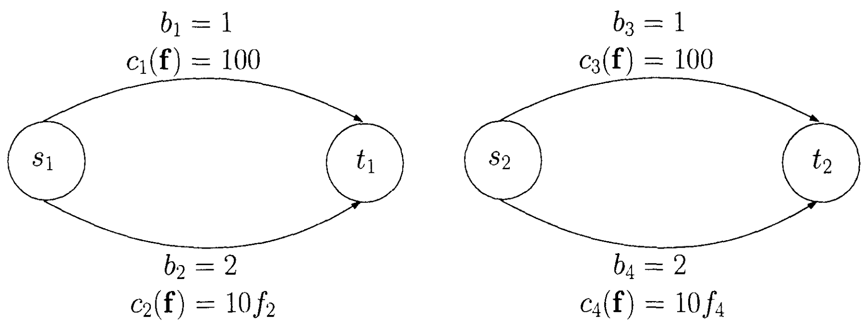

Figure 5-5: A simple example of $RNDP_{UO}(D)$

Because the flow associated with each O-D pair in this example does not affect the costs on the arcs of the other O-D pair, we may decouple the two O-D pairs in evaluating $\tau_{UO}(D;\mathbf{y})$. In this way, in this example, it is relatively easy to compute $\tau_{UO}(D;\mathbf{y})$ exactly, $\forall \mathbf{y} \in Y$. Despite its simplicity, the example allows us to clarify the algorithm's features.

The purpose of the lower bounds computed in the course of the branch and bound algorithm are relatively self-explanatory. Nonetheless, to illustrate their use, suppose, at some point in the execution of the branch and bound algorithm on this example, we reach the full solution $\mathbf{y}^* = (0,1,1,0)$. Suppose that, at this point in the execution, this full solution is the best full solution found so far, with $\tau_{UO}(D;\mathbf{y}^*) = 2340$. Furthermore, suppose the algorithm next considers either adding partial solution $\boldsymbol{\pi} = (0,1,0,-1)$ to Ω or branching on $\boldsymbol{\pi}$. For the yes-completion $\mathbf{y}^1 = (0,1,0,1)$ of $\boldsymbol{\pi}$, the algorithm computes $\tau_{SO}(D;\mathbf{y}^1) = 2890$. This value provides a lower bound on $\tau_{UO}(D;\mathbf{y})$, for all connectivity feasible descendants $\mathbf{y}$ of $\boldsymbol{\pi}$. In light of this bound, even if $\mathbf{y}^* = (0,1,1,0)$ is a suboptimal solution (in fact, it is suboptimal), any feasible descendant of $\boldsymbol{\pi}$ is even worse than $\mathbf{y}^*$. Thus, the algorithm excludes $\boldsymbol{\pi}$ from further consideration.

Next, let us illustrate the purpose of the upper bounds computed in the course of the algorithm. Suppose, in this example, we add partial solution $\boldsymbol{\pi} = (1,1,1,-1)$ to the list Ω of "candidate" partial solutions requiring further attention. One may verify that $\tau_{SO}(D;\mathbf{y}^0) = 2050$, where $\mathbf{y}^0 = (1,1,1,0)$, i.e., is the no-completion of $\boldsymbol{\pi}$. In addition, note that $\mathbf{c}(\mathbf{f})$ is a separable vector function. Accordingly, it follows from Roughgarden and Tardos [143, 144] and Roughgarden [141] that the price of anarchy, for any subnetwork of the one pictured in Figure 5-5, is bounded by 4/3. Thus,

$$\tau_{UO}(D; \mathbf{y}) \;\leq\; \frac{4}{3}\tau_{SO}\left(D; \mathbf{y}^0\right) \;=\; 2733\frac{1}{3}, \qquad \forall \mathbf{y} \in S(\boldsymbol{\pi}) \cap Y.$$

That is, any full solution $\mathbf{y}$ that descends from partial solution $\boldsymbol{\pi}$ yields a worst-case total system cost under UO routing of at most $2733\frac{1}{3}$. Suppose that the branch and bound algorithm considers adding $\boldsymbol{\omega} = (1, -1, 0, -1)$ to the candidate list Ω or branching on $\boldsymbol{\omega}$, if it is already in Ω. Let $\mathbf{y}^1 = (1, 1, 0, 1)$, i.e., the yes-completion of $\boldsymbol{\omega}$. Since

$$\tau_{UO}(D; \mathbf{y}) \;\geq\; \tau_{SO}\left(D; \mathbf{y}^1\right) \;=\; 2800, \qquad \forall \mathbf{y} \in S(\boldsymbol{\omega}) \cap Y,$$

all descendants of $\boldsymbol{\omega}$ yield higher worst-case costs than any descendant of $\boldsymbol{\pi}$. Consequently, the branch of the tree emanating from $\boldsymbol{\omega}$ may be excluded from further consideration; all of its descendants are clearly suboptimal. In this way, our branch and bound algorithm prescriptively uses the concept of the price of anarchy in an effort to eliminate suboptimal branches of the branch and bound tree and to thereby save on computational effort.

Fully carrying out the branch and bound algorithm on this example, and always branching on the partial solution in Ω with corresponding minimal lower bound, we correctly identify the exact optimal solution $\mathbf{y}^* = (1, 1, 1, 0)$, with corresponding $\tau_{UO}\left(D; \mathbf{y}^*\right) = 2140$. In the course of executing the algorithm, we reach only six of the sixteen total leaf nodes in the branch and bound tree. In addition, the algorithm requires that we evaluate $\tau_{UO}(D; \mathbf{y})$ at only four of these six leaf nodes, as the other two are infeasible.

Heuristic Solution of $RNDP_{UO}(D)$ via Branch and Bound Algorithm

As would any exact algorithm for solving $RNDP_{SO}(D)$ and $RNDP_{UO}(D)$, the exact version of the branch and bound algorithm we propose requires computation of $\tau_{SO}(D; \mathbf{y})$ and $\tau_{UO}(D; \mathbf{y})$, respectively, for various feasible solutions considered in the course of the algorithm. When $\tau_{SO}(D; \mathbf{y})$ can be computed exactly, the algorithm provides an exact solution method for $RNDP_{SO}(D)$. Recall that, when the arc costs are monotone and affine, and D is a bounded polyhedron, $\tau_{SO}(D;\mathbf{y})$ may be computed as the maximum value of $\zeta_{SO}(\mathbf{y},\mathbf{d})$ over the extreme points of D. We summarize the major features of the exact branch and bound algorithm for $RNDP_{SO}(D)$ in Figure 5-6. Furthermore, when both $\tau_{SO}(D;\mathbf{y})$ and $\tau_{UO}(D;\mathbf{y})$ can be computed exactly, the branch and bound algorithm provides an exact solution method for $RNDP_{UO}(D)$.

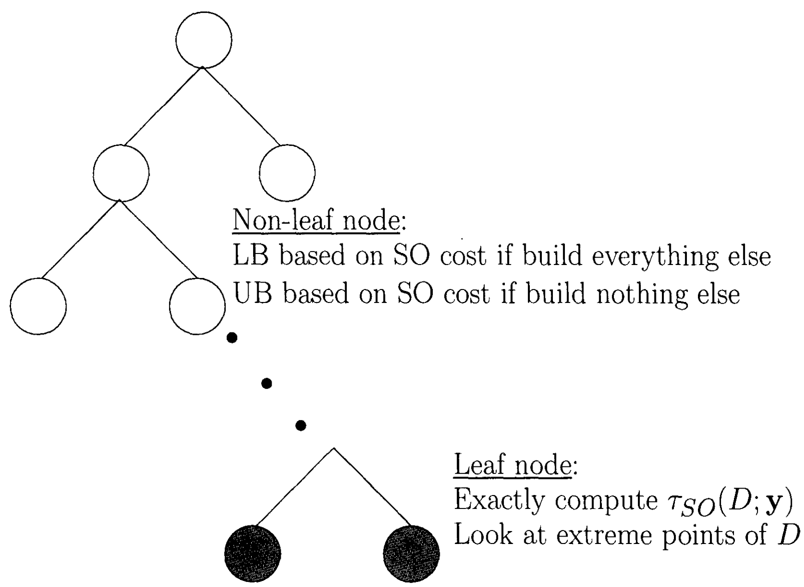

Figure 5-6: Features of the exact branch and bound algorithm for $RNDP_{SO}(D)$

As we discussed in Section 5.3, it is difficult to exactly evaluate $\tau_{UO}(D;\mathbf{y})$, even when D is a polyhedron and the arc cost functions are monotone and affine. In such cases, as proposed in Section 5.3.3, one can compute, via Lagrangian relaxation, an upper bound on $\tau_{UO}(D;\mathbf{y})$ in place of its true value. Accordingly, one can thereby use the following variation on the proposed branch and bound algorithm to heuristically solve $RNDP_{UO}(D)$.

In particular, for $\mathbf{y} \in Y$, let $L^*(D,\mathbf{y})$ denote the upper bound on $\tau_{UO}(D;\mathbf{y})$ obtained by Lagrangian relaxation. Recall from the proof of Theorem 5.4.1 that $\alpha \cdot \tau_{SO}(D;\mathbf{y})$ is also an upper bound on $\tau_{UO}(D;\mathbf{y})$. Consider a variation of the branch and bound algorithm for $RNDP_{UO}(D)$ in which, in Steps 3.b and 4.b of the algorithm, one uses the minimum of the two aforementioned upper bounds, i.e.,

$$\min\Big\{L^*(D,\mathbf{y}), \alpha \cdot \tau_{SO}\left(D;\mathbf{y}\right)\Big\}, \tag{5.25}$$

in place of $\tau_{UO}(D;\mathbf{y})$ itself. This variation of the algorithm solves for the feasible design solution for which upper bound (5.25) on $\tau_{UO}(D;\mathbf{y})$ is minimal. In Figure 5-7, we illustrate the main features of this version of the algorithm.

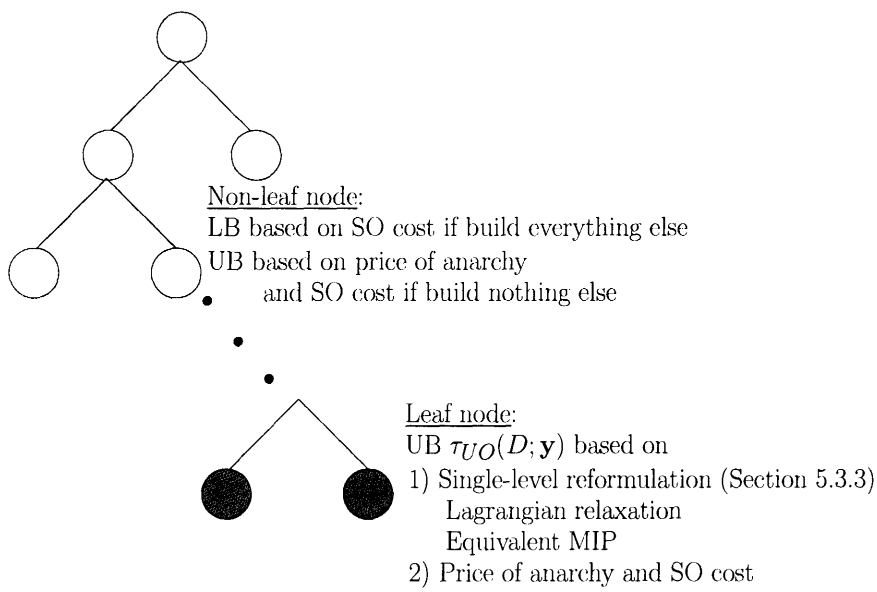

Figure 5-7: Features of the heuristic branch and bound algorithm for $RNDP_{UO}(D)$

For this heuristic solution $\mathbf{y}^*$, we would like the algorithm to output a range of values to which $\tau_{UO}\left(D;\mathbf{y}^*\right)$, though it cannot be computed exactly, is guaranteed to belong. We have established a heuristic version of the branch and bound algorithm that already computes the upper bound of such a range. We now describe how to augment the algorithm so that it also computes a lower bound on $\tau_{UO}\left(D;\mathbf{y}^*\right)$.

To begin, recall from the proof of Theorem 5.4.1 that $\tau_{UO}\left(D;\mathbf{y}\right) \geq \tau_{SO}\left(D;\mathbf{y}\right)$, $\forall \mathbf{y} \in Y$. For a second valid lower bound, recall that, for any given $\mathbf{y} \in Y$, $\tau_{UO}(D;\mathbf{y})$ is itself an MPEC. We established in Section 5.3.3 that, when D is a polyhedron and the arc cost functions are monotone and affine, this MPEC is equivalent to a nonconvex QCQP. Any local optimum of this QCQP thus yields a second valid lower bound on $\tau_{UO}(D; \mathbf{y}^*)$. Accordingly, we add the following final step to the heuristic version of the branch and bound algorithm for $RNDP_{UO}(D)$. At termination, the algorithm computes a lower bound on $\tau_{UO}(D; \mathbf{y}^*)$ by computing the maximum of $\tau_{SO}(D; \mathbf{y}^*)$ and a locally optimal objective value of the aforementioned QCQP. Note that, when $\mathbf{y}^*$ is the yes-completion of its parent partial solution, it inherits the lower bound $\tau_{SO}(D; \mathbf{y}^*)$ from this parent.

In this way, the heuristic version of our branch and bound algorithm not only identifies the feasible design solution $\mathbf{y}^*$ for which upper bound (5.25) on $\tau_{UO}(D; \mathbf{y})$ is minimal, but also determines a range of values to which $\tau_{UO}(D; \mathbf{y}^*)$ is guaranteed to belong.

Other Implementation Issues

Let us now justify and comment on measures built into the branch and bound algorithm to save on computational effort. Note that calls to $\tau_\rho(D; \mathbf{y})$ are not limited to steps in which the algorithm reaches a leaf node. Indeed, at a non-leaf node $\boldsymbol{\pi}$, in order to evaluate $LB_\rho(S(\boldsymbol{\pi}))$ and $UB_\rho(S(\boldsymbol{\pi}))$, the algorithm must compute $\tau_{SO}(D; \mathbf{y})$, where $\mathbf{y}$ is the appropriate completion of $\boldsymbol{\pi}$. Fortunately, however, the algorithm need not compute these bounds at every branch step. Indeed, consider a partial solution $\boldsymbol{\pi}$ with $\pi_a = -1$, $a \in A$. Consider, as given in equation (5.24), the children nodes $\boldsymbol{\pi}^1$ and $\boldsymbol{\pi}^0$ induced by branching on arc a. Since

$$LB_\rho\left(S\left(\boldsymbol{\pi}^1\right)\right) \quad = \quad LB_\rho\left(S\left(\boldsymbol{\pi}\right)\right),$$

$\boldsymbol{\pi}^1$ inherits its lower bound from $\boldsymbol{\pi}$. Similarly, if the no-completion of $\boldsymbol{\pi}$ belongs to Y, then so does the no-completion of $\boldsymbol{\pi}^0$, and furthermore,

$$UB_\rho\left(S\left(\boldsymbol{\pi}^0\right)\right) \quad = \quad UB_\rho\left(S\left(\boldsymbol{\pi}\right)\right).$$

In this way, since $\boldsymbol{\pi}^0$ inherits its upper bound from $\boldsymbol{\pi}$, the algorithm must update UB^* only when it considers child node $\boldsymbol{\pi}^1$.

In addition to evaluating $\tau_\rho(D; \mathbf{y})$, the branch and bound algorithm must assess whether partial solutions have any descendants that are feasible with respect to the construction budget and with respect to the the connectedness of the O-D pairs. The former assessment can be made by computing the total construction cost of the no-completion of the partial solution in question. The latter may be made by performing breadth first search, or a similar method, on the graph induced by the yes-completion of the partial solution. Again, the algorithm need not make these assessments at every branch step. The reason is that, if $\exists \mathbf{y} \in S(\boldsymbol{\pi})$ such that $\mathbf{b}'\mathbf{y} \leq B$, then $\exists \mathbf{y} \in S(\boldsymbol{\pi}^0)$ such that $\mathbf{b}'\mathbf{y} \leq B$. As such, the algorithm need only check budget feasibility when considering child node $\boldsymbol{\pi}^1$. Similarly, if $\exists \mathbf{y} \in S(\boldsymbol{\pi})$ such that $\mathbf{y} \in Y$, then $\exists \mathbf{y} \in S(\boldsymbol{\pi}^1)$ such that $\mathbf{y} \in Y$. Consequently, the algorithm need only verify connectivity when considering child node $\boldsymbol{\pi}^0$.

Lastly, let us discuss the breadth of NDP instances to which our branch and bound algorithm may be applied. The version of the algorithm we propose addresses robust (or nominal) NDPs in which the network designer is planning the edge set "from scratch," rather than adding to an existing arc set. One can easily modify the algorithm to address the latter scenario. Specifically, in this setting, one should initialize Ω to be the singleton partial solution $\boldsymbol{\pi}$ such that the following holds. $\pi_a = 1$ if a is a pre-existing arc, and $\pi_a = -1$, otherwise. In addition, in this context, if the no-completion of $\boldsymbol{\pi}$ is connectivity feasible, one should initialize UB^* differently. In particular, rather than initializing $UB^* := \infty$, one should instead initialize UB^* to $\alpha \cdot UB_\rho(S(\boldsymbol{\pi}))$.

5.6 A Single-level QCLP Reformulation of $RNDP_{UO}(D)$

In the previous section, we proposed a branch and bound algorithm for solving the robust NDP. Even in the most favorable circumstances, whether in the context of $RNDP_{SO}(D)$ or $RNDP_{UO}(D)$, carrying out this branch and bound algorithm requires significant computational effort. As such, we wish to identify classes of instances of these problems for which better solution methods exist. In Section 5.3.2, with Theorems 5.3.2 and 5.3.3, we identified conditions under which $RNDP_{UO}(D)$ and $RNDP_{SO}(D)$, respectively, reduce to nominal problems. These instances involve networks with only a single O-D pair and arbitrary network topology.

In this section, we examine $RNDP_{UO}(D)$ settings involving multiple O-D pairs, each connected by a set of parallel arcs, with a vector of separable, affine, and strictly monotone arc cost functions. For two reasons, this class of instances of $RNDP_{UO}(D)$ admits a simpler solution approach relative to the branch and bound algorithm. First, in this setting, we are able to parametrically solve, in closed form, for the unique UO flow solution, as a function of $\mathbf{y}$ and $\mathbf{d}$. As a result, we are able to solve for $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ in closed form. Recall that, in contrast, such a closed-form solution is not usually attainable for the general case. We derive these closed-form solutions in Section 5.6.1. Second, we are able to characterize a necessary and sufficient condition for all realizations of $\tilde{\mathbf{d}}$ from D to yield UO flow solutions in which, $\forall \mathbf{y} \in Y$, all built arcs are used. We establish this necessary and sufficient condition in Section 5.6.2.

As we discuss in Section 5.6.3, these two properties allow us to formulate $RNDP_{UO}(D)$ as a min-max problem over $\mathbf{y} \in Y$ and $\tilde{\mathbf{d}} \in D$, respectively. Exploiting convexity properties relating to the inner maximization problem, and using duality techniques from robust optimization, we can rewrite the min-max as a single-level optimization problem over polynomials. Specifically, this problem is a quadratically constrained linear optimization problem (QCLP). In part because it captures the binary integer constraints on $\mathbf{y}$, this QCLP is a nonconvex problem. However, in light of recent advances in (nonconvex) optimization over polynomials (see Lasserre [85, 86], Henrion and Lasserre [70], Parrilo and Sturmfels [122], Parrilo [121], and Prajna et al. [131]), for small instances of $RNDP_{UO}(D)$, this solution approach may be attractive.

5.6.1 Solving the NEP in Closed Form under Parallel Arcs

Consider a graph $G(V, A)$ with a set W of O-D pairs, each connected by a collection of parallel arcs. Equivalently, since $P = A$ in this setting, we may refer to these parallel arcs as paths. For the sake of clarity of notation, let us continue the discussion using the path, rather than the arc, terminology. Let $n_w = |P_w|$, that is, the number of paths (equivalently, the number of arcs), connecting O-D pair $w \in W$ in $G(V, A)$.

For the purpose of the discussion in this section of the thesis, let us introduce notation, for the vector of flow variables and vector of path (i.e., arc) cost functions, whose numbering convention differs slightly from the one we use in the rest of this chapter. In particular, let the vector of path flows be denoted by $\mathbf{F} = (\mathbf{F}^1, \ldots, \mathbf{F}^{|W|}) \in \mathbb{R}^{|P|}$, where $\mathbf{F}^w \in \mathbb{R}^{n_w}$ and F_p^w, $p \in \{1, \ldots, n_w\}$, denotes the flow on the p^{th} path connecting O-D pair w. Similarly, let $\mathbf{C}(\mathbf{F}) = (\mathbf{C}^1(\mathbf{F}), \ldots, \mathbf{C}^{|W|}(\mathbf{F})) \in \mathbb{R}^{|P|}$, where $\mathbf{C}^w(\mathbf{F}) \in \mathbb{R}^{n_w}$ and $C_p^w(\mathbf{F})$, $p \in \{1, \ldots, n_w\}$, denotes the cost induced by flow $\mathbf{F}$ on the p^{th} path connecting O-D pair w. Let us similarly define slightly different notation for the vector of binary design decisions. Let $\mathbf{y} = (\mathbf{y}^1, \ldots, \mathbf{y}^{|W|}) \in \{0,1\}^{|P|}$, where $y_p^w = 1$ denotes that the p^{th} path connecting O-D pair w is built, and $y_p^w = 0$ denotes that it is not built.

Suppose that the path cost functions are given as follows.

$$C_p^w(\mathbf{F}) \;=\; g_p^w F_p^w + h_p^w, \qquad w \in W;\; p \in \{1, \ldots, n_w\}, \tag{5.26}$$

where $g_p^w > 0$ and $h_p^w \geq 0$, $\forall w \in W$, $\forall p \in \{1, \ldots, n_w\}$. This setting is pictured in Figure 5-8. Thus, the path cost functions are affine, separable, and strictly monotone. Accordingly,

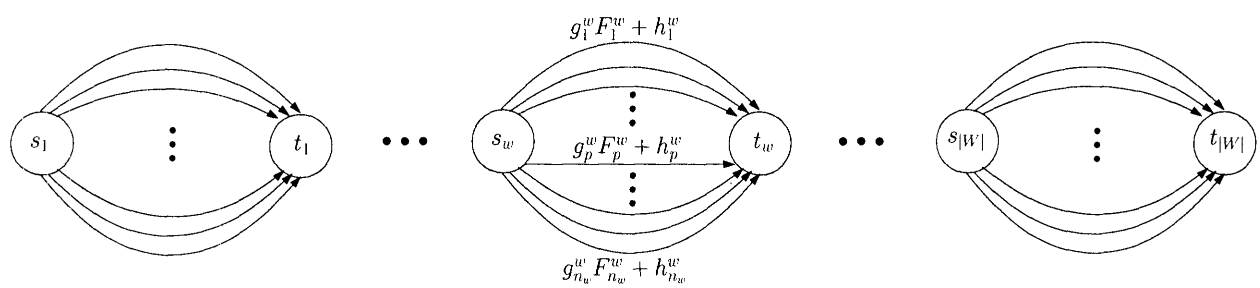

Figure 5-8: An NDP with parallel paths and separable and affine path cost functions

$\forall \mathbf{y} \in Y$, $\forall \mathbf{d} \in \mathbb{R}_+^{|W|}$, there is a unique equilibrium flow solution (Theorem 5.2.2). Without loss of generality, assume that, $\forall w \in W$,

$$h_1^w \;\leq\; h_2^w \;\leq\; \cdots \;\leq\; h_{n_w}^w. \tag{5.27}$$

In the following proposition, we solve for the unique equilibrium, in closed form, as a function of $\mathbf{y}$ and $\mathbf{d}$. Our result generalizes the exact equilibration algorithm of Dafermos and Sparrow [39] to the setting involving a variable arc set that depends on $\mathbf{y}$. We first establish some notation to be used in the proposition. For $w \in W$, $p \in \{1, \ldots, n_w - 1\}$, let

$$I_p^w(\mathbf{y}) = \left[\sum_{j=1}^{p} \frac{\left(h_p^w - h_j^w\right) y_j^w}{g_j^w}, \sum_{j=1}^{p+1} \frac{\left(h_{p+1}^w - h_j^w\right) y_j^w}{g_j^w} \right],$$

and let

$$I_{n_w}^w(\mathbf{y}) = \left[\sum_{j=1}^{n_w} \frac{\left(h_{n_w}^w - h_j^w\right) y_j^w}{g_j^w}, \infty \right).$$

For $p \in \{1, \ldots, n_w - 1\}$, note that the two endpoints of interval $I_p^w(\mathbf{y})$ are equal whenever $h_{p+1}^w = h_p^w$. Therefore, we include both endpoints in the interval in order to ensure that the interval is well defined when this equality holds. Note that, $\forall w \in W$, $\forall p \in \{1, \ldots, n_w - 1\}$,

$$\sum_{j=1}^{p+1} \frac{\left(h_{p+1}^w - h_j^w\right) y_j^w}{g_j^w} = \sum_{j=1}^{p} \frac{\left(h_{p+1}^w - h_j^w\right) y_j^w}{g_j^w} \geq \sum_{j=1}^{p} \frac{\left(h_p^w - h_j^w\right) y_j^w}{g_j^w}.$$

Thus, $\forall w \in W$, $\forall p \in \{1, \ldots, n_w\}$, $I_p^w(\mathbf{y})$ is indeed well defined.

Proposition 5.6.1. *Consider a graph $G(V, A)$ with a set W of O-D pairs, each connected by parallel paths. Suppose that the vector of path cost functions $\mathbf{C}(\mathbf{F})$ is given by (5.26) and satisfies the ordering in (5.27). Consider any $\mathbf{y} \in Y$. Let $\mathbf{d} \in \mathbb{R}_+^{|W|}$, with $\mathbf{d} > 0$. Suppose $d_w \in I_p^w(\mathbf{y})$, $w \in W$, where $p \in \{1, \ldots, n_w\}$. Let $\mathbf{F}(\mathbf{y}, \mathbf{d})$ denote the unique UO flow solution induced by $\mathbf{y}$ and $\mathbf{d}$. Then*

$$F_q^w(\mathbf{y}, \mathbf{d}) = \begin{cases} y_q^w \left[\dfrac{\sum_{j=1}^{p} \frac{\left(h_j^w - h_q^w\right) y_j^w}{g_j^w} + d_w}{\sum_{j=1}^{p} \frac{g_q^w y_j^w}{g_j^w}} \right], & q = 1, \ldots, p \\ \\ 0, & q = p+1, \ldots, n_w. \end{cases} \tag{5.28}$$

In addition, $\forall w \in W$,

$$\lambda_w(\mathbf{y}, \mathbf{d}) \;=\; \frac{\sum_{j=1}^{p} \frac{h_j^w y_j^w}{g_j^w} + d_w}{\sum_{j=1}^{p} \frac{y_j^w}{g_j^w}}.$$

Finally, the total cost in the system at this unique UO flow solution is

$$\zeta_{UO}(\mathbf{y}, \mathbf{d}) \;=\; \sum_{w \in W} d_w \lambda_w(\mathbf{y}, \mathbf{d}) \;=\; \sum_{w \in W} \frac{\sum_{j=1}^{p} \frac{h_j^w y_j^w d_w}{g_j^w} + d_w^2}{\sum_{j=1}^{p} \frac{y_j^w}{g_j^w}}.$$

Proof. By strict monotonicity of $\mathbf{C}$ and Theorem 5.2.2, $\forall \mathbf{y} \in Y$, and $\forall \mathbf{d} \in \mathbb{R}_+^{|W|}$, there exists a unique UO flow solution.

Consider $F(\mathbf{y}, \mathbf{d})$, as given by equation (5.28). Let us first establish that the expressions in (5.28) are well defined. Recall that $\forall w \in W$, $\forall j \in \{1, \ldots, n_w\}$, $g_j^w > 0$, where g_j^w is as in (5.26). From $\mathbf{y} \in Y$, it follows that $\forall w \in W$, $\exists j \in \{1, \ldots, n_w\}$ such that $y_j^w = 1$. Accordingly, the denominators of the fractions appearing in $\mathbf{F}$ are all nonzero, and $F(\mathbf{y}, \mathbf{d})$, as given by equation (5.28), is well defined.

We must next show that $F(\mathbf{y}, \mathbf{d}) \in K_P(\mathbf{y}, \mathbf{d})$, i.e., that it is a feasible vector of path flows with respect to $\mathbf{y}$ and $\mathbf{d}$. From (5.28) itself, we may conclude that $y_q^w = 0$ implies $F_q^w = 0$. $\forall w \in W$, let $p \in \{1, \ldots, n_w\}$ denote the index such that $d_w \in I_p^w(\mathbf{y})$. From the definition of $I_p^w(\mathbf{y})$, we have

$$d_w \;\geq\; \sum_{j=1}^{p} \frac{\left(h_p^w - h_j^w\right) y_j^w}{g_j^w}.$$

Adding $\sum_{j=1}^{p} \frac{\left(h_j^w - h_q^w\right) y_j^w}{g_j^w}$ to both sides, and using the fact that $h_q^w \leq h_p^w$, $\forall q \in \{1, \ldots, p\}$, we obtain

$$\sum_{j=1}^{p} \frac{\left(h_j^w - h_q^w\right) y_j^w}{g_j^w} + d_w \;\geq\; \sum_{j=1}^{p} \frac{\left(h_p^w - h_q^w\right) y_j^w}{g_j^w} \;\geq\; 0.$$

This nonnegativity relation implies that $\mathbf{F} \geq \mathbf{0}$. Lastly, $\forall w \in W$,

$$\sum_{q=1}^{n_w} F_q^w(\mathbf{y}, \mathbf{d}) \;=\; \sum_{q=1}^{p} y_q^w \left[\frac{\sum_{j=1}^{p} \frac{(h_j^w - h_q^w) y_j^w}{g_j^w} + d_w}{\sum_{j=1}^{p} \frac{g_q^w y_j^w}{g_j^w}} \right] \;=\; d_w.$$

Thus, $F(\mathbf{y}, \mathbf{d}) \in K_P(\mathbf{y}, \mathbf{d})$.

Having established that $F(\mathbf{y}, \mathbf{d})$ is a feasible vector of path flows, let us show that it satisfies Wardrop's Principle, given in Definition 5.2.4. $\forall w \in W$, $\forall q \in \{1, \ldots, p\}$ such that $y_q^w = 1$ (thereby implying $F_q^w(\mathbf{y}, \mathbf{d}) \geq 0$), we have

$$g_q^w F_q^w(\mathbf{y}, \mathbf{d}) + h_q^w \;=\; g_q^w \left[\frac{\sum_{j=1}^{p} \frac{(h_j^w - h_q^w) y_j^w}{g_j^w} + d_w}{\sum_{j=1}^{p} \frac{g_q^w y_j^w}{g_j^w}} \right] + h_q^w \;=\; \frac{\sum_{j=1}^{p} \frac{h_j^w y_j^w}{g_j^w} + d_w}{\sum_{j=1}^{p} \frac{y_j^w}{g_j^w}}. \tag{5.29}$$

Furthermore, $d_w \in I_p^w(\mathbf{y})$ implies that

$$d_w \;\leq\; \sum_{j=1}^{p+1} \frac{\left(h_{p+1}^w - h_j^w\right) y_j^w}{g_j^w} \;=\; \sum_{j=1}^{p} \frac{\left(h_{p+1}^w - h_j^w\right) y_j^w}{g_j^w}.$$

This inequality, together with equation (5.29), implies that

$$g_q^w F_q^w(\mathbf{y}, \mathbf{d}) + h_q^w \;\leq\; h_{p+1}^w \;\leq\; h_{p+2}^w \;\leq\; \cdots \;\leq\; h_{n_w}^w.$$

Thus, $F(\mathbf{y}, \mathbf{d})$ satisfies Wardrop's Principle. □

Note that, $\forall w \in W$, if d_w exactly equals an endpoint of an interval, then d_w belongs to $I_p^w(\mathbf{y})$ for at least two distinct $p \in \{1, \ldots, n_w\}$. Suppose that $p_1 \in \{1, \ldots, n_w - 1\}$ and $p_2 \in \{2, \ldots, n_w\}$, such that $p_2 \geq p_1 + 1$ and such that $d_w \in I_{p_1}^w(\mathbf{y})$ and $d_w \in I_{p_2}^w(\mathbf{y})$. $d_w \in I_{p_1}^w(\mathbf{y})$ implies

$$d_w \;\leq\; \sum_{j=1}^{p_1+1} \frac{\left(h_{p_1+1}^w - h_j^w\right) y_j^w}{g_j^w},$$

while $d_w \in I^w_{p_2}(\mathbf{y})$ implies

$$d_w \;\geq\; \sum_{j=1}^{\ell} \frac{\left(h^w_\ell - h^w_j\right) y^w_j}{g^w_j} \;\geq\; \sum_{j=1}^{p_1+1} \frac{\left(h^w_{p_1+1} - h^w_j\right) y^w_j}{g^w_j}, \qquad \ell = p_1+1, p_1+2, \ldots, p_2.$$

Therefore, $d_w \in I^w_{p_1}(\mathbf{y})$ and $d_w \in I^w_{p_2}(\mathbf{y})$ imply that

$$\begin{aligned} d_w &= \sum_{j=1}^{\ell} \frac{\left(h^w_\ell - h^w_j\right) y^w_j}{g^w_j}, && \ell = p_1+1, p_1+2, \ldots, p_2, \\ d_w &\in I^w_\ell(\mathbf{y}), && \ell = p_1, p_1+1, \ldots, p_2. \end{aligned} \tag{5.30}$$

Membership condition (5.30) may arise with $p_2 \geq p_1 + 1$ if

$$h^w_\ell \;=\; h^w_{p_1+1}, \qquad \ell = p_1+1, p_1+2, \ldots, p_2, \tag{5.31}$$

or if $h^w_\ell > h^w_{\ell-1}$ and $\ell \in \{p_1+2, \ldots, p_2\}$ imply that $y^w_j = 0$, $\forall j \in \{1, 2, \ldots, \ell-1\}$. In the latter case, since $\ell - 1 \geq p_1 + 1$, it follows that $d_w = 0$, thereby violating the assumption in Proposition 5.6.1 that $\mathbf{d} > \mathbf{0}$. Therefore, within the context of Proposition 5.6.1, membership condition (5.30) implies (5.31).

Under (5.30) and (5.31), setting $p = \ell$ in formula (5.28) yields the same value of $\mathbf{F}(\mathbf{y}, \mathbf{d})$, $\forall \ell \in \{p_1, p_1+1, \ldots, p_2\}$. Indeed, $\forall q \in \{p_1+1, \ldots, p_2\}$, $F^w_q(\mathbf{y}, \mathbf{d}) = 0$, since, $\forall \ell \in \{p_1 + 1, p_1+2, \ldots, p_2\}$,

$$\sum_{j=1}^{\ell} \frac{\left(h^w_j - h^w_q\right) y^w_j}{g^w_j} + d_w \;=\; \sum_{j=1}^{\ell} \frac{\left(h^w_j - h^w_\ell\right) y^w_j}{g^w_j} + d_w \;=\; 0.$$

In addition, $\forall q \in \{1, \ldots, p_1\}$, $\forall \ell \in \{p_1, p_1+1, \ldots, p_2\}$,

$$\frac{\sum_{j=1}^{\ell} \frac{(h^w_j - h^w_q) y^w_j}{g^w_j} + d_w}{\sum_{j=1}^{\ell} \frac{g^w_q y^w_j}{g^w_j}} \;=\; \frac{\sum_{j=1}^{\ell} \frac{(h^w_j - h^w_q) y^w_j}{g^w_j} + \sum_{j=1}^{\ell} \frac{(h^w_{p_1+1} - h^w_j) y^w_j}{g^w_j}}{\sum_{j=1}^{\ell} \frac{g^w_q y^w_j}{g^w_j}} \;=\; \frac{h^w_{p_1+1} - h^w_q}{g^w_q},$$

from which we obtain that

$$F_q^w(\mathbf{y}, \mathbf{d}) \;=\; y_q^w \left[\frac{h_{p_1+1}^w - h_q^w}{g_q^w}\right].$$

The next corollary follows directly from Proposition 5.6.1 and algebraic manipulation. We therefore omit its proof.

Corollary 5.6.1. *Consider a graph $G(V, A)$ with a set W of O-D pairs, each connected by parallel paths. Suppose that the vector of path cost functions $\mathbf{C}(\mathbf{F})$ is given by (5.26) and satisfies the ordering in (5.27). Consider any $\mathbf{y} \in Y$ and any $\mathbf{d} \in \mathbb{R}_+^{|W|}$, with $\mathbf{d} > 0$. Then, $\forall w \in W$,*

$$\lambda_w(\mathbf{y}, \mathbf{d}) \;=\; \min_{p \in \{1,\ldots,n_w\}} \frac{\sum_{j=1}^{p} \frac{h_j^w y_j^w}{g_j^w} + d_w}{\sum_{j=1}^{p} \frac{y_j^w}{g_j^w}}.$$

That is, for $\mathbf{y} \in Y$ fixed, $\lambda_w(\mathbf{y}, \mathbf{d})$ is a piecewise-linear, concave, and strictly increasing function of d_w.

Note that Proposition 5.6.1 and Corollary 5.6.1 both demonstrate that, as one would expect in this setting, $\lambda_w(\mathbf{y}, \mathbf{d})$ does not depend on $d_{w'}$, for $w' \in W$ with $w' \neq w$.

5.6.2 Heavy Traffic Conditions

Consider the closed-form formulae, given in Proposition 5.6.1, for the unique $\mathbf{F}(\mathbf{y}, \mathbf{d}) \in UO_P(\mathbf{y}, \mathbf{d})$, for the corresponding $\lambda_{UO}(\mathbf{y}, \mathbf{d})$, and for $\zeta_{UO}(\mathbf{y}, \mathbf{d})$, where $\mathbf{y} \in Y$ and $\mathbf{d} > \mathbf{0}$ are arbitrary. Recall that our goal in deriving these quantities is to reformulate

$$RNDP_{UO}(D) \;:\; \min_{\mathbf{y}} \left\{ \max_{\tilde{\mathbf{d}}} \left\{ \zeta_{UO}\left(\mathbf{y}, \tilde{\mathbf{d}}\right) \;\middle|\; \tilde{\mathbf{d}} \in D \right\} \;\middle|\; \mathbf{y} \in Y,\; \mathbf{b}'\mathbf{y} \leq B \right\}$$

in order to facilitate this problem's solution. The formulae given in Proposition 5.6.1 are conditioned on d_w belonging to $I_p^w(\mathbf{y})$, where $p \in \{1, \ldots, n_w\}$. In this way, p, and therefore the formulae themselves, implicitly depend on $\mathbf{y}$ and $\mathbf{d}$. Because of this implicit dependence, the formulae may be of little help in reframing $RNDP_{UO}(D)$ as a single-level optimization problem.

Corollary 5.6.1 allows us to rewrite $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ in closed form as an explicit function of $\mathbf{y}$ and $\mathbf{d}$.

$$\zeta_{UO}(\mathbf{y}, \mathbf{d}) = \sum_{w \in W} d_w \lambda_w(\mathbf{y}, \mathbf{d}) = \sum_{w \in W} d_w \min_{p \in \{1, \ldots, n_w\}} \frac{\sum_{j=1}^{p} \frac{h_j^w y_j^w}{g_j^w} + d_w}{\sum_{j=1}^{p} \frac{y_j^w}{g_j^w}}.$$

However, again, this formula for $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ is too complicated to be useful in manipulating $RNDP_{UO}(D)$ in the way that we would like.

In this section, we characterize "heavy traffic" conditions in which, $\forall \mathbf{d} \in D$, $\forall \mathbf{y} \in Y$, $\forall w \in W$, d_w is guaranteed to belong to $I_{n_w}^w(\mathbf{y})$. In such settings, the dependence on $\mathbf{y}$ and $\mathbf{d}$ of the expressions for the UO flow, for $\lambda_w(\mathbf{y}, \mathbf{d})$, and for $\zeta_{UO}(\mathbf{y}, \mathbf{d})$ automatically becomes explicit, without complicated reformulation. Essentially, this condition ensures that, for any design decision $\mathbf{y}$, and for any demand realization $\mathbf{d}$ in D, $\mathbf{y}$ and $\mathbf{d}$ induce an equilibrium flow solution in which all built paths (equivalently, in this setting, arcs) carry positive flow.

Corollary 5.6.2. *Consider a graph $G(V, A)$ with a set W of O-D pairs, each connected by parallel paths. Suppose that the vector of path cost functions $\mathbf{C}(\mathbf{F})$ is given by (5.26), and satisfies the ordering in (5.27). Let $\mathbf{F}(\mathbf{y}, \mathbf{d})$ denote the unique equilibrium flow solution corresponding to design vector $\mathbf{y} \in Y$ and demand vector $\mathbf{d} \in \mathbb{R}_+^{|W|}$. Consider a set D of demand vectors. The following two conditions are equivalent.*

Condition 1) $\forall \mathbf{d} \in D$, $\forall w \in W$, $d_w > \sum_{p=1}^{n_w} \frac{h_{n_w}^w - h_p^w}{g_p^w}$.

Condition 2) $\forall \mathbf{y} \in Y$, $\forall \mathbf{d} \in D$, $\forall w \in W$, $\forall p \in \{1, \ldots, n_w\}$, $y_p^w = 1$ *implies* $F_p^w(\mathbf{y}, \mathbf{d}) > 0$.

Proof. Suppose that Condition 1 holds. Then, $\forall \mathbf{d} \in D$, $\forall w \in W$,

$$d_w > \sum_{p=1}^{n_w} \frac{h_{n_w}^w - h_p^w}{g_p^w} \geq \sum_{p=1}^{n_w} \frac{\left(h_{n_w}^w - h_p^w\right) y_p^w}{g_p^w}, \tag{5.32}$$

implying that, $d_w \in I^w_{n_w}(\mathbf{y})$, $\forall \mathbf{y} \in Y$. Therefore, by Proposition 5.6.1, $\forall w \in W$, $\forall q \in \{1, \ldots, n_w\}$,

$$F^w_q(\mathbf{y}, \mathbf{d}) \;=\; y^w_q \left[\frac{\sum_{p=1}^{n_w} \frac{(h^w_p - h^w_q) y^w_p}{g^w_p} + d_w}{\sum_{p=1}^{n_w} \frac{g^w_q y^w_p}{g^w_p}} \right].$$

From inequality (5.32), we have

$$\sum_{p=1}^{n_w} \frac{\left(h^w_p - h^w_q\right) y^w_p}{g^w_p} + d_w \;>\; \sum_{p=1}^{n_w} \frac{\left(h^w_{n_w} - h^w_q\right) y^w_p}{g^w_p} \;\geq\; 0.$$

Thus, $y^w_p = 1$ implies $F^w_p(\mathbf{y}, \mathbf{d}) > 0$.

For the reverse direction, suppose that Condition 2 holds. Consider $\mathbf{y} = \mathbf{e}$, the vector of all ones. Consider any $\mathbf{d} \in D$. Condition 2 states that $F^w_p(\mathbf{y}, \mathbf{d}) > 0$, $\forall w \in W$, $\forall p \in \{1, \ldots, n_w\}$. Since $\mathbf{F}$ is a UO flow solution, it follows that, $\forall p \in \{1, \ldots, n_w\}$,

$$g^w_p F^w_p(\mathbf{y}, \mathbf{d}) + h^w_p \;=\; \cdots \;=\; g^w_{n_w} F^w_{n_w}(\mathbf{y}, \mathbf{d}) + h^w_{n_w} \;>\; h^w_{n_w},$$

where the inequality follows from the facts that $g^w_{n_w} > 0$. Dividing through by g^w_p and summing over $p \in \{1, \ldots, n_w\}$, we obtain

$$d_w \;=\; \sum_{p=1}^{n_w} F^w_p(\mathbf{y}, \mathbf{d}) \;>\; \sum_{p=1}^{n_w} \frac{h^w_{n_w} - h^w_p}{g^w_p}.$$

Thus, Condition 1 holds. □

5.6.3 A QCLP Reformulation of $RNDP_{UO}(D)$

In this section, we prove that, under a particular form of uncertainty set and under the heavy traffic conditions characterized in Section 5.6.2, $RNDP_{UO}(D)$ may be reformulated as a single-level QCLP.

Theorem 5.6.1. *Consider a graph $G(V, A)$ with a set W of O-D pairs, each connected by* *parallel paths. Let the vector of path cost functions* $\mathbf{C}(\mathbf{F})$ *be given by (5.26) and satisfy the ordering in (5.27). Let* $\mathbf{b} = \left(\mathbf{b}^1, \ldots, \mathbf{b}^{|W|}\right)$*, where* b_p^w *denotes the cost of constructing the* p^{th} *path (i.e., arc) connecting the* w^{th} *O-D pair. Let*

$$D \;=\; \left\{ \tilde{\mathbf{d}} \in \mathbb{R}^{|W|} \;\middle|\; \tilde{d}_w \in \left[\check{d}_w, \check{d}_w + \hat{d}_w\right], w \in W; \; \sum_{w \in W} \frac{\tilde{d}_w - \check{d}_w}{\hat{d}_w} \le \Gamma \right\},$$

where $\check{\mathbf{d}} > \mathbf{0}$*,* $\hat{\mathbf{d}} \ge \mathbf{0}$*, and* $\Gamma \ge 0$ *are given parameters. Furthermore, suppose that* D *satisfies Condition 1 of Corollary 5.6.2. Then,* $RNDP_{UO}(D)$ *is equivalent to the following QCLP.*

$$\begin{aligned}
\min_{\substack{\mathbf{y} \in \mathbb{R}^{|P|} \\ r \in \mathbb{R} \\ \mathbf{v}, \mathbf{z} \in \mathbb{R}^{|W|}}} \quad & \Gamma r + \sum_{w \in W} (v_w + z_w) \\
s.t. \quad & (v_w + r) \sum_{p=1}^{n_w} \frac{y_p^w}{g_p^w} \;\ge\; \hat{d}_w \sum_{p=1}^{n_w} \frac{h_p^w y_p^w}{g_p^w} + 2\check{d}_w \hat{d}_w + \hat{d}_w^2, && w \in W \\
& z_w \sum_{p=1}^{n_w} \frac{y_p^w}{g_p^w} \;=\; \check{d}_w \sum_{p=1}^{n_w} \frac{h_p^w y_p^w}{g_p^w} + \check{d}_w^2, && w \in W \qquad (5.33) \\
& \mathbf{b}'\mathbf{y} \;\le\; B \\
& \mathbf{e}'\mathbf{y}^w \;\ge\; 1, && w \in W \\
& y_p^w \left(1 - y_p^w\right) \;=\; 0, && w \in W; \; p \in \{1, \ldots, n_w\} \\
& \mathbf{v} \;\ge\; \mathbf{0} \\
& r \;\ge\; 0,
\end{aligned}$$

where $\mathbf{e}$ *is the appropriately dimensioned vector of all ones.*

Proof. To begin, let us rewrite

$$D \;=\; \left\{ \tilde{\mathbf{d}} \in \mathbb{R}^{|W|} \;\middle|\; \exists \mathbf{u} \in \mathbb{R}^{|W|} \text{ s.t. } \tilde{d}_w = \check{d}_w + u_w \hat{d}_w, \; w \in W; \; \mathbf{0} \le \mathbf{u} \le \mathbf{e}, \; \sum_{w \in W} u_w \le \Gamma \right\},$$

where $\mathbf{e} \in \mathbb{R}^{|W|}$ is the vector of all ones. For any $\mathbf{0} \le \mathbf{u} \le \mathbf{e}$, let $\tilde{\mathbf{d}} \in D$ be given by

$\tilde{d}_w = \check{d}_w + u_w \hat{d}_w$. From Condition 1 of Corollary 5.6.2, and from Proposition 5.6.1,

$$\begin{aligned}
\zeta_{UO}(\mathbf{y}, \tilde{\mathbf{d}}) &= \sum_{w \in W} \frac{\sum_{p=1}^{n_w} \frac{h_p^w y_p^w \tilde{d}_w}{g_p^w} + \tilde{d}_w^2}{\sum_{p=1}^{n_w} \frac{y_p^w}{g_p^w}} \\
&= \sum_{w \in W} \left[\frac{1}{\sum_{p=1}^{n_w} \frac{y_p^w}{g_p^w}} \right] \left[\sum_{p=1}^{n_w} \frac{h_p^w y_p^w}{g_p^w} \left(\check{d}_w + \hat{d}_w u_w \right) + \check{d}_w^2 + 2\check{d}_w \hat{d}_w u_w + \hat{d}_w^2 u_w^2 \right].
\end{aligned}$$

Therefore,

$$\begin{aligned}
\tau_{UO}(D; \mathbf{y}) &= \sum_{w \in W} \left[\frac{1}{\sum_{p=1}^{n_w} \frac{y_p^w}{g_p^w}} \right] \left[\sum_{p=1}^{n_w} \frac{h_p^w y_p^w}{g_p^w} \check{d}_w + \check{d}_w^2 \right] \\
&+ \max_{\mathbf{u}} \quad \sum_{w \in W} \left[\frac{1}{\sum_{p=1}^{n_w} \frac{y_p^w}{g_p^w}} \right] \left[\sum_{p=1}^{n_w} \frac{h_p^w y_p^w}{g_p^w} \hat{d}_w u_w + 2\check{d}_w \hat{d}_w u_w + \hat{d}_w^2 u_w^2 \right] \qquad (5.34) \\
&\qquad \text{s.t.} \quad \mathbf{0} \le \mathbf{u} \le \mathbf{e} \\
&\qquad\qquad \sum_{w \in W} u_w \le \Gamma.
\end{aligned}$$

Let $\mathbf{u}^{**}$ denote an optimal solution of LCQP (5.34).

In addition, let $\mathbf{u}^*$ denote an optimal solution of the following feasible and bounded linear program (LP) in $\mathbf{u}$. Without loss of generality, we may assume that $\mathbf{u}^* \in \{0, 1\}^{|W|}$.

$$\begin{aligned}
\max_{\mathbf{u}} \quad & \sum_{w \in W} \left[\frac{1}{\sum_{p=1}^{n_w} \frac{y_p^w}{g_p^w}} \right] \left[\sum_{p=1}^{n_w} \frac{h_p^w y_p^w}{g_p^w} \hat{d}_w + 2\check{d}_w \hat{d}_w + \hat{d}_w^2 \right] u_w \\
\text{s.t.} \quad & \mathbf{0} \le \mathbf{u} \le \mathbf{e} \qquad (5.35) \\
& \sum_{w \in W} u_w \le \Gamma.
\end{aligned}$$

The following system establishes that the optimal values of LCQP (5.34) and LP (5.35) are equivalent. The equality in this system holds by the fact that $\mathbf{u}^* \in \{0, 1\}^{|W|}$. The first inequality follows from the fact that $\mathbf{u}^*$ is a feasible, but not necessarily optimal solution of LCQP (5.34). From $\mathbf{0} \le \mathbf{u}^{**} \le \mathbf{e}$, the second inequality holds. Finally, the third inequality follows from the fact that $\mathbf{u}^{**}$ is a feasible, but not necessarily optimal solution of LP (5.35).

$$\sum_{w\in W}\left[\frac{1}{\sum_{p=1}^{n_w}\frac{y_p^w}{g_p^w}}\right]\left[\sum_{p=1}^{n_w}\frac{h_p^w y_p^w}{g_p^w}\hat{d}_w+2\check{d}_w\hat{d}_w+\hat{d}_w^2\right]u_w^*$$

$$=\sum_{w\in W}\left[\frac{1}{\sum_{p=1}^{n_w}\frac{y_p^w}{g_p^w}}\right]\left[\sum_{p=1}^{n_w}\frac{h_p^w y_p^w}{g_p^w}\hat{d}_w u_w^*+2\check{d}_w\hat{d}_w u_w^*+\hat{d}_w^2\left(u_w^*\right)^2\right]$$

$$\leq\sum_{w\in W}\left[\frac{1}{\sum_{p=1}^{n_w}\frac{y_p^w}{g_p^w}}\right]\left[\sum_{p=1}^{n_w}\frac{h_p^w y_p^w}{g_p^w}\hat{d}_w u_w^{**}+2\check{d}_w\hat{d}_w u_w^{**}+\hat{d}_w^2\left(u_w^{**}\right)^2\right]$$

$$\leq\sum_{w\in W}\left[\frac{1}{\sum_{p=1}^{n_w}\frac{y_p^w}{g_p^w}}\right]\left[\sum_{p=1}^{n_w}\frac{h_p^w y_p^w}{g_p^w}\hat{d}_w+2\check{d}_w\hat{d}_w+\hat{d}_w^2\right]u_w^{**}$$

$$\leq\sum_{w\in W}\left[\frac{1}{\sum_{p=1}^{n_w}\frac{y_p^w}{g_p^w}}\right]\left[\sum_{p=1}^{n_w}\frac{h_p^w y_p^w}{g_p^w}\hat{d}_w+2\check{d}_w\hat{d}_w+\hat{d}_w^2\right]u_w^*.$$

Since LP (5.35) is feasible and bounded, its dual is also feasible and bounded and has an optimal value equal to that of LP (5.35). This dual LP is given by

$$\begin{aligned}\min_{r\in\mathbb{R},\ \mathbf{v}\in\mathbb{R}^{|W|}}\quad & \Gamma r+\sum_{w\in W}v_w\\ \text{s.t.}\quad & v_w+r\ \geq\ \left[\frac{1}{\sum_{p=1}^{n_w}\frac{y_p^w}{g_p^w}}\right]\left[\sum_{p=1}^{n_w}\frac{h_p^w y_p^w}{g_p^w}\hat{d}_w+2\check{d}_w\hat{d}_w+\hat{d}_w^2\right], \qquad w\in W\\ & \mathbf{v}\ \geq\ \mathbf{0}\\ & r\ \geq\ 0.\end{aligned}$$

Since $\mathbf{y}\in\{0,1\}^{|P|}$ and $\mathbf{g}>\mathbf{0}$, we may multiply the first constraint through by $\sum_{p=1}^{n_w}\frac{y_p^w}{g_p^w}$. Furthermore, letting

$$z_w\ =\ \frac{1}{\sum_{p=1}^{n_w}\frac{y_p^w}{g_p^w}}\left[\check{d}_w\sum_{p=1}^{n_w}\frac{h_p^w y_p^w}{g_p^w}+\check{d}_w^2\right],\qquad w\in W,$$

we obtain that $RNDP_{UO}(D)$ is equivalent to (5.33). □

5.7 Conclusions

In this chapter of the thesis, we propose a novel, distribution-free, robust optimization model of the binary choice, arc construction NDP, under demand uncertainty, congestion effects, and either UO or SO routing. We offer methods for solving the resulting robust NDP. In particular, we propose a branch and bound algorithm for exactly or heuristically solving instances under SO or UO routing, respectively. Whereas the price of anarchy has previously been used only in a descriptive, rather than a prescriptive, manner, our branch and bound algorithm constructively uses this measure of inefficiency. Moreover, we prove that the robust NDP under SO routing gives a price-of-anarchy-approximate solution to the robust NDP under UO routing. In addition, we present conditions under which the robust NDP reduces to a less computationally demanding problem. We give two such sufficient conditions. Under the first, the robust NDP reduces to a nominal counterpart, and under the second, it is equivalent to a single-level quadratic optimization problem. Furthermore, we observe counterintuitive behavior, not yet noted in the literature, of costs at equilibrium with respect to changes in traffic demands on the network. The examples we present are analogous to Braess' Paradox [27] and illustrate that an increase in traffic demands on a network may yield a strict decrease in the costs at equilibrium. Finally, we establish convexity and monotonicity properties of functions relating to the worst-case performance of a given network design decision. These properties motivate the solution methods we propose and underscore the relative levels of difficulty among the SO and UO versions of the robust NDP and their nominal counterparts.

Chapter 6

Conclusions

In this thesis, we offer a novel, distribution-free, robust optimization approach to three classes of equilibrium-related problems.

In the first part of the thesis, we explore the nominal variational inequality (VI) problem. Interestingly, although the data in the problem are regarded as being known with certainty, the VI is in fact a special instance of a robust constraint. Exploiting this insight, we use duality-based proof techniques from the robust optimization literature in order to reformulate the VI over a polyhedron as a single-level optimization problem. This reformulation applies even if the associated cost function has an asymmetric Jacobian matrix. Moreover, in contrast to other reformulation approaches in the VI literature which produce problems that are at most once differentiable, our reformulation yields an optimization problem that is many-times continuously differentiable. We give sufficient conditions for the convexity of this reformulation and thereby identify a class of VIs which may be solved using widely-available and commercial-grade convex optimization software. We prove that monotone affine (and possibly asymmetric) VIs are special instances of this class.

In the second part of the thesis, we consider an equilibrium-related problem that does involve data uncertainty, namely the finite game with incomplete information. Using the robust optimization paradigm, we propose a distribution-free model of incomplete-information games, in which the players use a robust optimization approach to contend with payoff uncertainty. Our "robust game" model relaxes the assumptions of Harsanyi's Bayesian game model, and provides an alternative, distribution-free equilibrium concept, for which, in contrast to *ex post* equilibria, existence is guaranteed. We show that computation of "robust-optimization equilibria" is analogous to that of Nash equilibria of complete-information games. Namely, for arbitrary robust finite games with bounded polyhedral payoff uncertainty sets, we provide a formulation of the set of robust-optimization equilibria as the dimension-reducing, component-wise projection of the solution set of a system of multilinear equalities and inequalities. We suggest that sample solutions of such systems can be approximately computed using a pseudo-Newton method applied to an appropriate penalty function. To demonstrate the practicality of this computational approach, we present numerical results of implementation. For a special class of robust finite games, we show that the set of robust-optimization equilibria is equivalent to the more easily computable Nash equilibria set of a complete-information game, with the same number of players and the same action spaces. Finally, we compare properties of robust finite games with those of the corresponding Bayesian games. Our results cover incomplete-information games without private information as well as those involving potentially private information.

Finally, in the third part of the thesis, we consider an alternate game-theoretical perspective on data uncertainty, namely, that of a mechanism designer. Specifically, we propose a novel, robust optimization model of the binary choice, arc construction network design problem (NDP), under demand uncertainty, congestion effects, and either system-optimal (SO) or user-optimal (UO) routing. We offer a branch and bound algorithm for solving the resulting robust NDP. This algorithm comprises the first constructive use of the price of anarchy concept, which has previously been employed only in a descriptive, rather than a prescriptive manner. Moreover, we prove that the optimal solution of the robust NDP under SO routing is a price-of-anarchy-approximate solution to the robust NDP under UO routing. In addition, we present conditions under which the robust NDP reduces to a less computationally demanding problem, either a nominal counterpart or a single-level quadratic optimization problem. Furthermore, we observe a novel traffic "paradox" relating to changes in costs at equilibrium with respect to changes in traffic demands on the network. The examples we present are analogous to Braess' Paradox [27] and illustrate that an increase in traffic demands on a network may yield a strict decrease in the costs at equilibrium. Finally, we establish convexity and monotonicity properties of functions relating to the worst-case performance of a given network design decision. These properties motivate the solution methods we propose and underscore the relative levels of difficulty among the SO and UO versions of the robust NDP and their nominal counterparts.

Thus, in this thesis, we extend the reach of the robust optimization paradigm to the field of game theory and to the nominal variational inequality problem. We thereby show that ideas from robust optimization are useful not only in settings characterized by data uncertainty, but also, interestingly, in contexts with known and certain data as well.

Appendix A

An Alternate Derivation of the VI Reformulation

In this appendix, we give an alternate proof, based on the KKT conditions for LP (3.6), of Theorem 3.2.1.

Proof. As mentioned in our first proof of Theorem 3.2.1, it is well known that $\mathbf{x}^* \in K$ solves $VI(K, \mathbf{F})$ iff $\mathbf{x}^*$ itself optimizes the LP (3.6) it induces. Consider K given, as before, by (3.4). Since the KKT conditions are necessary and sufficient for the optimality of an LP, $\mathbf{x}^*$ solves $VI(K, \mathbf{F})$ iff $\exists \boldsymbol{\lambda}^* \in \mathbb{R}^m,\ \boldsymbol{\mu}^* \in \mathbb{R}^n$ such that $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = (\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ satisfies the following KKT system corresponding to LP (3.6). This equivalence is also well known in the VI literature (see, e.g., Proposition 1.2.1 of Facchinei and Pang [47]).

$$\begin{aligned} \mathbf{F}(\mathbf{x}) - \mathbf{A}'\boldsymbol{\lambda} - \boldsymbol{\mu} &= \mathbf{0} \\ \mathbf{A}\mathbf{x} &= \mathbf{b} \\ \mathbf{x} &\geq \mathbf{0} \\ \boldsymbol{\mu} &\geq \mathbf{0} \\ \mu_j x_j &= 0, \qquad j \in \{1, \ldots, n\}. \end{aligned} \tag{A.1}$$

Suppose $\exists \boldsymbol{\lambda}^* \in \mathbb{R}^m,\ \boldsymbol{\mu}^* \in \mathbb{R}^n$ such that $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = (\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ satisfies KKT system (A.1). Then, from the first constraint and the fact that $\boldsymbol{\mu}^* \geq \mathbf{0}$, it follows that

$$\mathbf{F}(\mathbf{x}^*) \;\geq\; \mathbf{A}'\boldsymbol{\lambda}^*.$$

In addition, taking the inner product of $\mathbf{x}^*$ and the left-hand side of the first constraint, and using the other constraints in (A.1), one obtains $0 = \mathbf{F}(\mathbf{x}^*)' \mathbf{x}^* - \mathbf{b}^{*\prime}\boldsymbol{\lambda}^*$. Therefore, $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ satisfies (3.5).

For the reverse direction, suppose that $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ satisfies system (3.5). Let $\boldsymbol{\mu}^* \in \mathbb{R}^n$ be given by

$$\mu_j^* \;=\; \begin{cases} 0, & x_j^* > 0 \\ F_j(\mathbf{x}^*) - \mathbf{A}_j'\boldsymbol{\lambda}^*, & x_j^* = 0, \end{cases}$$

where $\mathbf{A}_j$ denotes the j^{th} column of the matrix $\mathbf{A}$. Clearly, $\boldsymbol{\mu}^* \geq \mathbf{0}$ and $\mu_j^* x_j^* = 0$, $\forall j \in \{1, \ldots, n\}$. Suppose $\exists j \in \{1, \ldots, n\}$ such that $x_j^* > 0$ and $F_j(\mathbf{x}^*) > A_j'\boldsymbol{\lambda}^*$. Then, $F(\mathbf{x}^*)' \mathbf{x}^* > \mathbf{b}'\boldsymbol{\lambda}^*$, thereby yielding a contradiction. Accordingly, $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ satisfies KKT system (A.1), i.e., $\mathbf{x}^*$ solves $VI(K, \mathbf{F})$. □

Appendix B

Comparison of Notions of Network Equilibrium

In this appendix, we compare Wardrop's concept of network equilibrium with an alternate definition, given by Dafermos and Sparrow [39]. Wardrop's equilibrium concept requires that all flow travels on minimum-cost paths. Dafermos' and Sparrow's notion of equilibrium requires that no fraction of the flow on one path may be rerouted to a different path in a way that strictly decreases the cost experienced by the rerouted flow.

Definition B.1.1 (Dafermos and Sparrow [39]). For a given $\mathbf{y} \in \{0,1\}^{|A|}$, and a given $\mathbf{d} \in \mathbb{R}_+^{|W|}$, $\mathbf{F} \in K_P(\mathbf{y}, \mathbf{d})$ is a **user-optimal**, or **equilibrium** vector of path flows for network $G(V, A(\mathbf{y}))$ with path costs $\mathbf{C}$ if, $\forall w \in W$, $\forall p_1, p_2 \in P_w(\mathbf{y})$, $\forall \delta \in (0, F_{p_1}]$,

$$f_{p_1} \; > \; 0 \quad \Longrightarrow \quad C_{p_1}(\mathbf{F}) \; \leq \; C_{p_2}(\mathbf{F}'),$$

where $\mathbf{F}' \in K_P(\mathbf{y}, \mathbf{d})$ is given by

$$F'_p \; = \; \begin{cases} F_p - \delta, & p = p_1, \\ F_p + \delta, & p = p_2, \\ F_p, & \text{otherwise.} \end{cases}$$

This equilibrium concept is sometimes referred to as a "Nash flow" (see, e.g. Roughgarden and Tardos [143, 144]).

We can more clearly contrast the two equilibrium concepts by recasting them in layman's terms. Essentially, Wardrop's equilibrium concept requires that no units of flow are "envious" of other units of flow associated with the same O-D pair (i.e., because the latter units experience a lower cost than the former). In contrast, a Nash flow solution requires that no flow has incentive to unilaterally deviate to a different path. In some settings, Wardrop's equilibrium concept and that of the Nash flow are equivalent.

Theorem B.1.1 (Dafermos and Sparrow [39]). *Consider a network $G(V, A)$ with continuous and monotone path cost functions $\mathbf{C}$. For a given $\mathbf{y} \in \{0,1\}^{|A|}$, and a given $\mathbf{d} \in \mathbb{R}_+^{|W|}$, $\mathbf{F} \in K_P(\mathbf{y}, \mathbf{d})$ satisfies Definition B.1.1, i.e., is a Nash flow, iff it is a Wardrop equilibrium, i.e., iff it satisfies Definition 5.2.4.*

In other settings, the two equilibrium concepts are not equivalent. For instance, recall Example 5.3.3 which we presented in Section 5.3.2. In this example, $\mathbf{f} = (f_1, f_2) = (1, 0)$ is a Wardrop equilibrium since $c_1(1,0) = c_2(1,0) = 10$. However, it is not a Nash flow, since some (in fact all) of the flow has incentive to unilaterally deviate to another path. Indeed, $c_2(1-\epsilon, \epsilon) < 10, \forall \epsilon \in (0, 1]$. Therefore, under arc costs that are not monotone with respect to the arc flows, a Wardrop equilibrium need not be a Nash flow; i.e., the incentive to deviate may occur without some units of flow "envying" other units of flow. Conversely, under discontinuous arc cost functions, a Nash flow need not be a Wardrop equilibrium; i.e., there may be no incentive to deviate, even when there is "envy." Consider the example given in Appendix B.1 of Roughgarden [142]. This example involves a graph like the one pictured in Example 5.3.3 in Section 5.3.2, but with arc cost functions

$$c_1(\mathbf{f}) = \begin{cases} 0, & f_1 \in \left[0, \frac{1}{3}\right], \\ 1, & f_1 > \frac{1}{3}, \end{cases} \qquad c_2(\mathbf{f}) = \begin{cases} \frac{1}{2}, & f_2 \in \left[0, \frac{1}{3}\right], \\ 1, & f_1 > \frac{1}{3}. \end{cases}$$

One may verify that $\mathbf{f} = (f_1, f_2) = \left(\frac{1}{3}, \frac{2}{3}\right)$ is a Nash flow. However, it is not a Wardrop equilibrium, since $c_1\left(\frac{1}{3}, \frac{2}{3}\right) < c_2\left(\frac{1}{3}, \frac{2}{3}\right)$.

Although this potential nonequivalence between the Wardrop equilibrium and Nash flow concepts is worth noting, it is not relevant in the majority of Chapter 5 of this thesis, in which our discussion relates to concepts of network equilibrium. Indeed, all vectors of arc cost functions we consider in Chapter 5 are continuous, and, with the exception of the vector of arc cost functions in Example 5.3.3, all are monotone with respect to the arc flows.

Bibliography

[1] Andrade, R., A. Lisser, N. Maculan, G. Plateau (2004): "Telecommunication Network Capacity Design for Uncertain Demand," *Computational Optimization and Applications*, 29: 127–46.

[2] Atamturk, A. and M. Zhang (2004): "Two-Stage Robust Network Flow and Design under Demand Uncertainty," Research Report BCOL 04.03.

[3] Auchmuty, G. (1989): "Variational Principles for Variational Inequalities," *Numerical Functional Analysis and Optimization* 10: 863–74.

[4] Balakrishnan, A. (1987): "LP Extreme Points and Cuts for the Fixed-Charge Network Design Problem," *Mathematical Programming*, 39: 263–84.

[5] Balakrishnan, A., T. Magnanti, P. Mirchandani (1997): "Network Design," *Annotated Bibliographies in Combinatorial Optimization*, Eds. M. Dell'Amico, F. Maffioli, S. Marello, Wiley, New York, 311–34.

[6] Balakrishnan, A., T. Magnanti, and R. Wong (1989): "A Dual-based Procedure for Large Scale Uncapacitated Network Design," *Operations Research*, 37: 716–40.

[7] Bazaraa, M. and C. Shetty (1993): *Nonlinear Programming: Theory and Algorithms*, 2nd ed., Wiley, New York.

[8] Beckmann, M., C. McGuire, and C. Winsten (1956): *Studies in the Economics of Transportation*, Yale University Press.

[9] Benders, J. (1962): "Partitioning Procedures for Solving Mixed Variable Programming Problems," *Num. Math.*, 4: 238–52.

[10] Ben-Tal, A., A. Goryashko, E. Guslitzer, and A. Nemirovski (2004): "Adjustable Robust Solutions of Uncertain Linear Programs," *Mathematical Programming*, 99: 351–376.

[11] Ben-Tal, A. and A. Nemirovski (1998): "Robust Convex Optimization," *Mathematics of Operations Research*, 23: 769–805.

[12] Ben-Tal, A. and A. Nemirovski (1999): "Robust Solutions of Uncertain Linear Programs," *Operations Research Letters*, 25: 1–13.

[13] Ben-Tal, A. and A. Nemirovski (2000): "Robust Solutions of Linear Programming Problems Contaminated with Uncertain Data," *Mathematical Programming, Series A*, 88:411–24.

[14] Ben-Tal, A. and A. Nemirovski (2001): *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*, MPS/SIAM Series on Optimization, SIAM, Philadelphia.

[15] Bertsekas, D. (1995): *Nonlinear Programming*, 2nd ed., Athena Scientific, Belmont.

[16] Bertsimas, D., D. Pachamanova, and M. Sim (2004): "Robust Linear Optimization under General Norms," *Operations Research Letters*, 32: 510–16.

[17] Bertsimas, D. and M. Sim (2003): "Robust Discrete Optimization and Network Flows," *Mathematical Programming, Series B*, 98: 49–71.

[18] Bertsimas, D. and M. Sim (2004): "The Price of Robustness," *Operations Research*, 52: 35–53.

[19] Bertsimas, D. and J. Tsitsiklis (1997): *Introduction to Linear Optimization*, Athena Scientific, Belmont.

[20] Bertsimas, D. and R. Weismantel (2005): *Optimization over Integers*, Dynamic Ideas, Belmont.

[21] Billheimer, J. and P. Gray, (1973): "Network Design with Fixed and Variable Cost Elements," *Transportation Science*, 7: 49–74.

[22] Boffey, T. and A. Hinxman (1979): "Solving for Optimal Network Problem," *European Journal of Operational Research*, 3: 386–93.

[23] Bohnenblust, H. and S. Karlin (1950): "On a Theorem of Ville," *Contributions to the Theory of Games*, eds. H. Kuhn and A. Tucker, Princeton UP, Vol. 1, 155–60.

[24] Borodin, A. and R. El-Yaniv (1998): *Online Computation and Competitive Analysis*, New York: Cambridge UP.

[25] Boyce, D., A. Farhi, and R. Weischedel (1973): "Optimal Network Problem: A Branch-and-Bound Algorithm," *Environ. Plan.*, 5: 519–33.

[26] Boyce, D. and J. Soberanes (1979): "Solutions to the Optimal Network Design Problem with Shipment Related to Transportation Cost," *Transportation Research*, 13B: 65–80.

[27] Braess, D. (1968): "Uber ein Paradoxon der Verkehrsplanung," *Unternehmensforschung*, 12: 258–68.

[28] Brouwer, L.E.J. (1912): "Uber Abbildung von Mannigfaltikeiten," *Mathematische Annalen*, 71: 97–115.

[29] Chau, C.K. and K.M. Sim (2003): "The Price of Anarchy for Non-Atomic Congestion Games with Symmetric Cost Maps and Elastic Demands," *Operations Research Letters*, 31: 327–34.

[30] Chekuri, C., G. Oriolo, M. Scutellà, F. Shepherd (2005): "Hardness of Robust Network Design," *Proceedings of INOC 2005, Book 2*, Eds. Gouveia and Mourao, Lisbon, pp. 455–461.

[31] Correa, J.R., A.S. Schulz, and N.E. Stier Moses (2004): "Selfish Routing in Capacitated Networks," *Mathematics of Operations Research*, 29: 961–976.

[32] Cottle, R. (1964): *Nonlinear Programs with Positively Bounded Jacobians*, Ph.D. dissertation, Department of Mathematics, UC Berkeley.

[33] Cottle, R., J. Pang, and R. Stone (1992): *The Linear Complementarity Problem*, Academic Press, Boston.

[34] Courant, R. (1943): "Variational Methods for the Solution of Problems with Equilibrium and Vibration," *Bull. Amer. Math. Soc.* 49: 1–23.

[35] Crainic, T. and M. Gendreau (2002): "Cooperative Parallel Tabu Search for Capacitated Network Design," *Journal of Heuristics*, 8: 601–27.

[36] Crainic, T., M. Gendreau, and J. Farvolden (2000): "A Simplex-Based Tabu Search Method for Capacitated Network Design," *INFORMS Journal on Computing* 12: 223–36.

[37] Crémer, J. and R. McLean (1985): "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica*, 53: 345–62.

[38] Dafermos, S. (1980): "Traffic Equilibria and Variational Inequalities," *Transportation Science*, 14: 42–54.

[39] Dafermos, S. and F. Sparrow (1969): "The Traffic Assignment problem for a General Network," *Journal of Research of the National Bureau of Standards*, Series B, 73B(2): 91–118.

[40] Datta, R.S. (2003): "Using Computer Algebra to Find Nash Equilibria," *Proceedings of the 2003 International Symposium on Symbolic and Algebraic Computation*, 74–79.

[41] Debreu, G. (1952): "A Social Equilibrium Existence Theorem," *Proceedings of the National Academy of Sciences, USA*, 38: 886–93.

[42] Dionne R. and M. Florian (1979): "Exact and Approximate Algorithms for Optimal Network Design," *Networks*, 9: 37–59.

[43] Dow, J. and S. Werlang (1994): "Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction," *Journal of Economic Theory*, 64: 305–24.

[44] Eaves, B.C. (1971): "On the Basic Theorem of Complementarity," *Mathematical Programming* 1: 68–75.

[45] El Ghaoui, L. and H. Lebret (1997): "Robust Solutions to Least-Squares Problems with Uncertain Data," *SIAM Journal of Matrix Analysis Applications*, 18: 1035–64.

[46] El Ghaoui, L., F. Oustry, and H. Lebret (1998): "Robust Solutions to Uncertain Semidefinite Programs," *SIAM Journal on Optimization*, 9: 33–52.

[47] Facchinei, F. and J. Pang (2003): *Finite-Dimensional Variational Inequalities and Complementarity Problems*, New York, Springer.

[48] Ferris, M. and C. Kanzow (2002): "Complementarity and Related Problems: A Survey," *Handbook of Applied Optimization*, Eds. P. Pardalos and M. Resende, Oxford UP, New York, pp. 514–530.

[49] Fiacco, A.V. and G.P. McCormick (1968): *Nonlinear Programming: Sequential Unconstrained Minimization Techniques*, New York: John Wiley & Sons.

[50] Florian, M. and D. Hearn (1995): "Network Equilibrium Models and Algorithms," *Network Routing*, eds. M. Ball, T. Magnanti, C. Monma, and G. Nemhauser, chapter 6, 485–550.

[51] Floudas, C. (2000): *Deterministic Global Optimization: Theory, Methods and Applications*, Kluwer, Boston.

[52] Floudas, C. and V. Visweswaran (1994): "Quadratic Optimization," *Handbook of Global Optimization*, eds. R. Horst and P. Pardalos, Kluwer: 217–270.

[53] Fudenberg, D. and D.K. Levine (1998): *The Theory of Learning in Games*, Cambridge: MIT Press.

[54] Fudenberg, D. and J. Tirole (1991): *Game Theory*, Cambridge: MIT Press.

[55] Fukushima, M. (1992): "Equivalent Differentiable Optimization Problems and Descent Methods for Asymmetric Variational Inequality Problems," *Mathematical Programming* 53: 99–110.

[56] Gendron, B., T. Crainic, A. Frangioni (1999): "Multicommodity Capacitated Network Design," *Telecommunications Network Planning*, Eds. P. Soriano and B. Sansr, Kluwer Academic, pp. 1–19.

[57] Geoffrion, A. and G. Graves (1974): "Multicommodity Distribution System Design by Benders Decomposition," *Management Science*, 5: 822–44.

[58] Gilboa, I. and D. Schmeidler (1989): "Maxmin Expected Utility with a Non-Unique Prior," *Journal of Mathematical Economics*, 18: 141–53.

[59] Goldberg, A., et al. (2002): "Competitive Auctions," submitted to *Games and Economic Behavior*.

[60] Govindan, S. and R. Wilson (2003): "A Global Newton Method to Compute Nash Equilibria," *Journal of Economic Theory*, 110: 65–86.

[61] Gutiérrez, G., P. Kouvelis, and A. Kurawarwala (1996): "A Robust Approach to Uncapacitated Network Design Problems," *European Journal of Operational Research*, 94: 362–76.

[62] Hall, M. (1978): "Properties of the Equilibrium State in Transportation Networks," *Transportation Science*, 12: 208–16.

[63] Hammond, J. (1984): *Solving Asymmetric Variational Inequality Problems and Systems of Equations with Generalized Nonlinear Programming Algorithms*, Ph.D. dissertation, Department of Mathematics, MIT, Cambridge.

[64] Harker, P. and J. Pang (1990): "Finite-dimensional Variational Inequality and Nonlinear Complementarity Problems: A Survey of Theory, Algorithms and Applications," *Mathematical Programming, Series B* 48: 161–220.

[65] Harker, P. and B. Xiao (1991): "A Polynomial-time Algorithm for Affine Variational Inequalities," *Applied Mathematics Letters* 4: 31–34.

[66] Harsanyi, J.C. (1967, 1968): "Games with Incomplete Information Played by 'Bayesian' Players, I - III" *Management Science*, 14: 159–82, 320–34, 486–502.

[67] Hartman, P. and G. Stampacchia (1966): "On Some Nonlinear Elliptic Differential Functional Equations," *Acta Mathematica* 115: 271–310.

[68] Haurie, A. and P. Marcotte (1985): "On the Relationship Between Nash-Cournot and Wardrop Equilibria," *Networks*, 15: 295–308.

[69] Hellstrand, J., T. Larson, and A. Migdalas (1992): "A Characterization of the Uncapacitated Network Design Polytope," *Operations Research Letters*, 12: 159–63.

[70] Henrion, D. and J. Lasserre (2003): "GloptiPoly : Global Optimization over Polynomials with MATLAB and SeDuMi," *ACM Trans. Math. Soft.*, 29: 165–194.

[71] Hoang, H. (1973): "A Computational Approach to the Selection of an Optimal Network," *Management Science*, 19: 488–98.

[72] Hoang, H. (1982): "Topological Optimization of Networks: A Nonlinear Mixed Integer Model Employing Generalized Benders Decomposition," *IEEE Trans. Automatic Control*, AC-27: 164–69.

[73] Hochbaum, D. and A. Segev (1989): "Analysis of a Flow Problem with Fixed Charges," *Networks*, 19: 291–312.

[74] Holmberg, K. and J. Hellstrand (1998): "Solving the Uncapacitated Network Design Problem by a Lagrangean Heuristic and Branch-and-Bound," *Operations Research*, 46: 247–59.

[75] Holmström, B. and R.B. Myerson (1983): "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, 51: 1799–1820.

[76] Holzman, R., N. Kfir-Dahav, D. Monderer, M. Tennenholtz (2004): "Bundling Equilibrium in Combinatorial Auctions," *Games and Economic Behavior*, 47: 104–23.

[77] Hyafil, N. and C. Boutilier (2004): "Regret Minimizing Equilibria and Mechanisms for Games with Strict Type Uncertainty," *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, July 2004, 268–277.

[78] Johnson, D.S., J.K. Lenstra, and A.H.G. Rinnooy Kan (1978): "The Complexity of the Network Design Problem," *Networks* 8: 279-85.

[79] Kakutani, S. (1941): "A Generalization of Brouwer's Fixed Point Theorem," *Duke Mathematical Journal*, 8: 457–59.

[80] Karamardian, S. (1972): "The Complementarity Problem," *Mathematical Programming* 2: 107–29.

[81] Klibanoff, P. (1993): "Uncertainty, Decision, and Normal Form Games," Manuscript, MIT.

[82] Knight, F. (1921): *Risk, Uncertainty and Profit*, Boston: Houghton Mifflin.

[83] Koutsoupias, E. and C.H. Papadimitriou (1999): "Worst-case Equilibria," *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pp. 404–13.

[84] Kumar, R. and N. Banerjee (2003): "Multicriteria Network Design Using Evolutionary Algorithm," *Lecture Notes in Computer Science*, 2724: 2179–90.

[85] Lasserre, J. (2001): "Global Optimization with Polynomials and the Problem of Moments," *SIAM J. Optimization*, 11: 796–817.

[86] Lasserre, J. (2002): "Semidefinite Programming vs. LP Relaxations for Polynomial Programming," *Mathematics of Operations Research*, 27: 347–360.

[87] LeBlanc, L. (1975): "An Algorithm for the Discrete Network Design Problem," *Transportation Science*, 9: 283–87.

[88] LeBlanc, L. and D. Boyce (1986): "Bilevel Programming Algorithm for Exact Solutions of the Networks Design Problem with User-Optimal Flows," *Transportation Research*, 20B: 259–65.

[89] Lemke, C. (1965): "Bimatrix Equilibrium Points and Mathematical Programming," *Management Science* 11: 681–89.

[90] Lemke, C. and J. Howson (1964): "Equilibrium Points of Bimatrix Games," *SIAM Journal of Applied Mathematics* 12: 413–23.

[91] Lin, H., T. Roughgarden, E. Tardos (2004): "A Stronger Bound on Braess's Paradox," *SODA 2004*, 333–334.

[92] Lions, J. and G. Stampacchia (1967): "Variational Inequalities," *Communications on Pure and Applied Mathematics* 20: 493–519.

[93] Lisser, A., A. Ouorou, J. Vial, and J. Gondzio (1999): "Capacity Planning under Uncertain Demand in Telecommunication Networks," Logilab Technical Report 99.13, Department of Management Studies, University of Geneva, Switzerland.

[94] Lo, K.C. (1996): "Equilibrium in Beliefs Under Uncertainty," *Journal of Economic Theory*, 71: 443–84.

[95] Los, M. and C. Lardinois (1980): "Combinatorial Programming, Statistical Optimization and the Optimal Transportation Network Problem," *Transportation Research*, 16B: 89–124.

[96] Luo, Z., J. Pang, and D. Ralph (1996): *Mathematical Programs with Equilibrium Constraints*, Cambridge UP, New York.

[97] Lüthi, H. (1985): "On the Solution of Variational Inequalities by the Ellipsoid Method," *Mathematics of Operations Research* 10: 515–22.

[98] Ma, T. (2002): *Banach-Hilbert Spaces, Vector Measures and Group Representations*, New Jersey: World Scientific.

[99] Magnanti, T. (1984): "Models and Algorithms for Predicting Urban Traffic Equilibria," *Transportation Planning Models*, ed. M. Florian, Amsterdam, 153-85.

[100] Magnanti, T., P. Mireault, and R. Wong (1986): "Tailoring Benders Decomposition for Uncapacitated Network Design," *Mathematical Programming Study*, 26: 112–154.

[101] Magnanti, T. and G. Perakis (1995): "A Unifying Geometric Solution Framework and Complexity Analysis for Variational Inequalities," *Mathematical Programming* 71: 327–51.

[102] Magnanti, T. and R. Wong (1984): "Network Design and Transportation Planning: Models and Algorithms," *Transportation Science*, 18: 1–55.

[103] Mancino, O. and G. Stampacchia (1972): "Convex Programming and Variational Inequalities," *Journal of Optimization Theory and Applications* 9: 3–23.

[104] Marcotte, P. (1983): "Network Optimization with Continuous Control Parameters," *Transportation Science*, 17: 181–97.

[105] Marcotte, P. (1986): "Network Design Problem with Congestion Effects: A Case of Bilevel Programming," *Mathematical Programming*, 34: 142–62.

[106] Marinacci, M. (2000): "Ambiguous Games," *Games and Economic Behavior*, 31: 191–219.

[107] McKelvey, R.D. and A.M. McLennan (1996): "Computation of Equilibria in Finite Games," *Handbook of Computational Economics*, Vol. I , ed. by H. Amman, D. Kendrick, and J. Rust, Elsevier, 87–142.

[108] McKelvey, R.D., A.M. McLennan, and T.L. Turocy (2004): "Gambit: Software Tools for Game Theory", Version 0.97.0.6. http://econweb.tamu.edu/gambit.

[109] Mertens, J.-F. and S. Zamir (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14: 1–29.

[110] Minoux, M. (1989): "Network Synthesis and Optimum Network Design Problems: Models, Solution Methods and Applications," *Networks*, 19: 313–360.

[111] Minoux, M. (2003): "Optimum Network Design Models and Algorithms in Transportation and Communication," *International Journal of Logistics: Research and Applications*, 6: 5–15.

[112] Morris, S. (1995): "The Common Prior Assumption in Economic Theory," *Economics and Philosophy*, 11: 227–53.

[113] Murchland, J. (1970): "Braess's Paradox of Traffic Flow," *Transportation Research*, 4: 391–94.

[114] Nagurney, A. (1993): *Network Economics: A Variational Inequality Approach*, Kluwer: Boston.

[115] Nagurney, A. (2000): "Congested Urban Transportation Networks and Emission Paradoxes," *Transportation Research Part D*, 5: 145–51.

[116] Nash, J.F., Jr. (1950): "Equilibrium Points in N-Person Games," *Proceedings of the National Academy of Sciences, USA*, 36: 48–49.

[117] Nash, J.F., Jr. (1951): "Non-Cooperative Games," *Annals of Mathematics*, 54: 286–95.

[118] Ordóñez, F. and J. Zhao (2004): "Robust Capacity Expansion of Network Flows," USC-ISE Working paper 2004-01.

[119] Osborne, M.J. and A. Rubinstein (1994): *A Course in Game Theory*, Cambridge: MIT Press.

[120] Papadimitriou, C. (2001): "Algorithms, Games, and the Internet," *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*, 749–53.

[121] Parrilo, P. (2005): "Exploiting Algebraic Structure in Sum of Squares Programs," *Positive Polynomials in Control: Lecture Notes in Control and Information Sciences*, 312: 181–194.

[122] Parrilo, P. and B. Sturmfels (2001): "Minimizing Polynomial Functions," *Algorithmic and Quantitative Real Algebraic Geometry: DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 60: 83–99.

[123] Patriksson, M. (1994): *The Traffic Assignment Problem: Models and Methods*, VSP, Utrecht, Netherlands.

[124] Patriksson, M. (1998): *Nonlinear Programming and Variational Inequality Problems: A Unified Approach*, vol. 23 of Applied Optimization, Kluwer, Dordrecht.

[125] Peng, J. (1997): "Equivalence of Variational Inequality Problems to Unconstrained Minimization," *Mathematical Programming* 78: 347–55.

[126] Perakis, G. (1992): *Geometric, Interior Point, and Classical Methods for Solving Finite Dimensional Variational Inequality Problems*, Ph.D. dissertation, Department of Applied Mathematics, Brown University, Providence, RI.

[127] Perakis, G. (2003): "The 'Price of Anarchy' under Nonlinear and Asymmetric Costs," to appear in *Mathematics of Operations Research*.

[128] Pigou, A. (1920): *The Economics of Welfare*, London, Macmillan.

[129] Plesnik, J. (1980): "The Complexity of Designing a Network with Minimum Diameter," *Networks*, 10: 33–45.

[130] Porter, R., E. Nudelman, and Y. Shoham (2004): "Simple Search Methods for Finding a Nash Equilibrium," *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI-2004).*

[131] Prajna, S., A. Papachristodoulou, and P. Parrilo (2005): "SOSTOOLS and its Control Applications," *Positive Polynomials in Control: Lecture Notes in Control and Information Sciences*, 312, 273–292.

[132] Randall, M., G. Graham McMahon, and S. Sugden (2002): "A Simulated Annealing Approach to Communication Network Design," *Journal of Combinatorial Optimization*, 6: 55–65.

[133] Rardin, R. and U. Choe (1979): "Tighter Relaxations of Fixed Charge Network Flow Problems," Technical Report J-79-18, Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta.

[134] Rardin, R. and L. Wolsey (1993): "Valid Inequalities and Projecting the Multicommodity Extended Formulation for Uncapacitated Fixed Charge Network Flow Problems," *European Journal of Operational Research*, 71: 95–109.

[135] Riis, M. and K. Andersen (2002): "Capacitated Network Design with Uncertain Demand," *INFORMS Journal on Computing*, 14: 247–260.

[136] Ronen, A. (2000): "Algorithms for rational agents," *Conference on Current Trends in Theory and Practice of Informatics*, 56–70.

[137] Rosenthal, R. (1973): "A Class of Games Possessing Pure Strategy Nash Equilibria," *Int. J. Game Theory*, 2: 65–67.

[138] Rosenthal, R. (1973): "The Network Equilibrium Problem in Integers," *Networks*, 3: 53–59.

[139] Rothengatter, W. (1979): "Application of Optimal Subset Selection to Problems of Design and Scheduling on Urban Transportation Networks," *Transportation Research*, 13B: 49–63.

[140] Roughgarden, T. (2001): "On the Severity of Braess's Paradox: Designing Networks for Selfish Users Is Hard," to appear in special FOCS 2001 issue of *Journal of Computer and System Sciences.*

[141] Roughgarden, T. (2002): "The Price of Anarchy is Independent of the Network Topology," *Proceedings 34th ACM Symposium on Theory of Computing*, 428–37.

[142] Roughgarden, T. (2002): *Selfish Routing*, Ph.D. dissertation, Department of Computer Science, Cornell University, Ithaca, NY.

[143] Roughgarden, T. and E. Tardos (2000): "How Bad is Selfish Routing?" *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science*, 93–102.

[144] Roughgarden, T. and E. Tardos (2002): "Bounding the Inefficiency of Equilibria in Nonatomic Congestion Games," *Journal of the ACM*, 49: 236–59.

[145] Scarf, H. (1967): "The Approximation of Fixed Points of a Continuous Mapping," *SIAM Journal on Applied Mathematics* 15: 1328–43.

[146] Scarf, H. with the collaboration of T. Hansen (1973): *The Computation of Economic Equilibria*, New Haven: Yale UP.

[147] Scott, A. (1969): "The Optimal Network Problem: Some Computational Procedures," *Transportation Research*, 3: 201–10.

[148] Smart, D. (1974): *Fixed Point Theorems*, Cambridge UP.

[149] Smith, M. (1979): "The Existence, Uniqueness and Stability of Traffic Equilibria," *Transportation Research*, 13B: 295–304.

[150] Soyster, A.L. (1973): "Convex Programming with Set-Inclusive Constraints and Applications to Inexact Linear Programming," *Operations Research*, 21: 1154–57.

[151] Stampacchia, G. (1964): "Formes Bilineares Coercives sur les Ensembles Convexes," *Comptes Rendus Academie Sciences Paris* 258: 4413–16.

[152] Stampacchia, G. (1968): "Variational Inequalities," *Theory and Applications of Monotone Operators: Proc. NATO Advanced Study Inst., Venice, 1968*, pp. 101–92.

[153] Sturmfels, B. (2002): *Solving Systems of Polynomial Equations*, Providence: American Mathematical Society.

[154] Sun, J. and G. Zhao (1998): "Global Linear and Local Quadratic Convergence of a Long-step Adaptive-mode Interior Point Method for Some Monotone Variational Inequality Problems," *SIAM Journal on Optimization* 8: 123–39.

[155] Sun, J. And G. Zhao (1998): "Quadratic Convergence of a Long-Step Interior-Point Method for Nonlinear Monotone Variational Inequality Problems," *Journal of Optimization Theory and Applications* 97: 471–91.

[156] Taji, K. and M. Fukushima (1996): "A New Merit Function and a Successive Quadratic Programming Algorithm for Variational Inequality Problems," *SIAM Journal on Optimization* 6: 704–13.

[157] Tseng, P. (1992): "Global Linear Convergence of a Path-following Algorithm for Some Monotone Variational Inequality Problems, *Journal of Optimization Theory and Applications* 75: 265–79.

[158] van der Laan, G. and A. Talman (1982): "On the Computation of Fixed Points in the Product Space of Unit Simplices and an Application to Noncooperative N Person Games," *Mathematics of Operations Research*, 7: 1–13.

[159] van der Laan, G., A. Talman, and L. van der Heyden (1987): "Simplicial Variable Dimension Algorithms for Solving the Nonlinear Complementarity Problem on a Product of Unit Simplices Using a General Labeling," *Mathematics of Operations Research* 12: 377–97.

[160] Verschelde, J. (1999): "PHCpack: A General-Purpose Solver for Polynomial Systems by Homotopy Continuation," Algorithm 795, *ACM Transactions on Mathematical Software*, 25: 251–76.

[161] von Neumann, J. and O. Morgenstern (1944): *The Theory of Games and Economic Behavior*, Princeton: Princeton UP.

[162] von Stengel, B. (2002): "Computing Equilibria for Two-Person Games," Chapter 45, *Handbook of Game theory with Economic Applications*, Vol. 3, eds. R.J. Aumann and S. Hart, Elsevier, Amsterdam, 1723–59.

[163] Waller, S.T. and A.K. Ziliaskopoulos (2001): "A Dynamic and Stochastic Approach to Network Design," *Journal of the Transportation Research Board*, 1771: 106–14.

[164] Wardrop, J.G. (1952): "Some Theoretical Aspects of Road Traffic Research," *Proceedings of the Institute of Civil Engineers*, Volume 1, Part II: 325–78.

[165] Wilson, R. (1987): "Game-Theoretic Approaches to Trading Processes," *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley. Chap. 2, pp. 33–77. Cambridge: Cambridge UP.

[166] Wong, R. (1980): "Worst-case Analysis of Network Design Problem Heuristics," *SIAM J. Algebraic and Discrete Methods*, 1: 51–63.

[167] Wright, S. (2001): "On Reduced Convex QP Formulations of Monotone LCPs," *Mathematical Programming*, Ser. A 90: 459–73.

[168] Wu, J. (1993): "Interior-point Algorithms for Some Monotone Variational Inequality Problems," Publication 879, Centre de Recherche sur les Transports, Université de Montréal.

[169] Wu, J. (1998): "Long-step Primal Path-following Algorithm for Monotone Variational Inequality Problems," *Journal of Optimization Theory and Applications* 99: 509–31.

[170] Wu, J., M. Florian, and P. Marcotte (1993): "A General Descent Framework for the Monotone Variational Inequality Problem," *Mathematical Programming* 61: 281–300.

[171] Yamashita, N. and M. Fukushima (1997): "Equivalent Unconstrained Minimization and Global Error Bounds for Variational Inequality Problems," *SIAM Journal on Control and Optimization* 35: 273–84.

[172] Yamashita, N., K. Taji, and M. Fukushima (1997): "Unconstrained Optimization Reformulations of Variational Inequality Problems," *Journal of Optimization Theory and Applications* 92: 439–56.

[173] Zhu, D. and P. Marcotte (1994): "An Extended Descent Framework for Variational Inequalities," *Journal of Optimization Theory and Applications* 80: 349–60.

[174] Zuhovickii, S., R. Poljak, and M. Primak (1969): "Two Methods of Search for Equilibrium Points of n-person Concave Games," *Soviet Mathematics Doklady* 10: 279–82.

[175] Zuhovickii, S., R. Poljak, and M. Primak (1969): "Numerical Methods of Finding Equilibrium Points of n-person Games," *Proceedings of the First Winter School of Mathematical Programming in Drogobych*, 93–130.

[176] Zuhovickii, S., R. Poljak, and M. Primak (1970): "On an n-person Concave Game and a Production Model," *Soviet Mathematics Doklady* 11: 522–26.

[177] Zuhovickii, S., R. Poljak, and M. Primak (1973): "Concave Equilibrium Points: Numerical Methods," *Matekon* 9: 10–30.